



**A Comparative Analysis of Proposed Mobility
Support Schemes For IP Multicast**

THESIS

Alexander Muller Jr., Captain, USAF

AFIT/GCS/ENG/00J-02

DEPARTMENT OF THE AIR FORCE
AIR UNIVERSITY

AIR FORCE INSTITUTE OF TECHNOLOGY

Wright-Patterson Air Force Base, Ohio

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

DTIC QUALITY INSPECTED 4

20010122 153

The views expressed in this thesis are those of the author and do not reflect the official policy or position of the United States Air Force, Department of Defense, or the U. S. Government

**A COMPARATIVE ANALYSIS OF PROPOSED MOBILITY SUPPORT
SCHEMES FOR IP MULTICAST**

THESIS

Presented to the faculty of the Graduate School of Engineering and Management

of the Air Force Institute of Technology

Air University

In Partial Fulfillment of the

Requirements for the Degree of

Master of Science in Computer Engineering

Alexander Muller Jr.

Captain, USAF

April 2000

Approved for public release; distribution unlimited

**A COMPARATIVE ANALYSIS OF PROPOSED MOBILITY SUPPORT
SCHEMES FOR IP MULTICAST**

THESIS

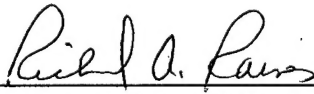
Presented to the faculty of the Graduate School of Engineering and Management
of the Air Force Institute of Technology
Air University

In Partial Fulfillment of the
Requirements for the Degree of
Master of Science in Computer Engineering

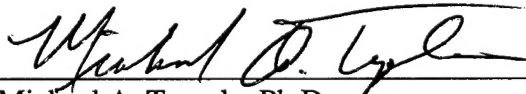
Alexander Muller Jr.

Captain, USAF

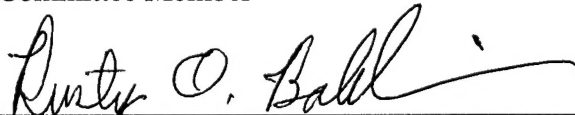
April 2000



Richard A. Raines, Ph.D., Major, USAF
Committee Chairman



Michael A. Temple, Ph.D.
Committee Member



Rusty O. Baldwin, Ph.D., Major, USAF
Committee Member

Approved for public release; distribution unlimited

Acknowledgements

There are several people to whom I owe a debt of gratitude for making it possible for me to complete this work. I would like first to thank my thesis advisor, Major Richard A. Raines, for his guidance and leadership throughout this endeavor. I would also like to thank my committee members, Doctor Michael A. Temple, and Major Rusty O. Baldwin for their advice, guidance, and technical expertise in our many research meetings.

I also owe many thanks to my fellow networks students with whom I spent many hours in the Hawkeye lab attempting to glean the nuances and subtleties of modeling with OPNET. Captains Brett Neville and Steve Conklin both provided fresh insights when I had racked my brain struggling with the coding mysteries of OPNET. Without these insights, I would have been unable to complete my models or this thesis.

The person to whom I owe the deepest and most heartfelt gratitude is my wonderful wife, Hyon-I. I could never put in words how grateful I am to Hyon-I for the effort that she put into taking care of our two small children, Emerald and Joshua. I doubt that I would ever be up to the Herculean task of spending 24 hours a day with two toddlers without a break for 18-months straight, but that's what Hyon-I did in order to allow me to devote my time and efforts to my studies. For this I am eternally grateful.

Finally, I would like to thank my Lord in heaven for giving Hyon-I the strength to endure, for helping Joshua overcome his illnesses, for making Emerald such a delight to come home to, and for giving me the strength to complete this work. To Him I owe everything.

Table of Contents

Acknowledgements	iv
Table of Contents	v
List of Figures	ix
List of Tables.....	x
List of Tables.....	x
Abstract	xi
Chapter 1: Introduction	1
1.1 Background	1
1.2 Research Goal	2
1.3 Research Motivation	2
1.4 Approach	4
1.5 Overview of Results	4
1.6 Summary	5
Chapter 2: Literature Review	6
2.1 Introduction	6
2.2 IP Multicast	6
2.2.1 Multicast Introduction	6
2.2.2 Host Extensions for IP Multicasting	9
2.2.3 Addressing.....	10
2.2.4 IGMP	13
2.2.5 Forwarding Algorithms	16
2.2.5.1 Flooding	18
2.2.5.2 Spanning Trees	18
2.2.5.3 Reverse Path Forwarding	19
2.2.5.4 Reverse Path Broadcast.....	19
2.2.5.5 Truncated Reverse Path Broadcast.....	20
2.2.5.6 Reverse Path Multicasting.....	20
2.2.5.7 Core-Based Trees	21
2.2.6 Current Implementations	23
2.2.6.1 Distance Vector Multicast Routing Protocol	24
2.2.6.2 Multicast OSPF	26
2.2.6.3 Protocol-Independent Multicast	26
2.3 IP in IP Tunneling	29
2.4 Mobile IP.....	31
2.4.1 Mobile IP Definitions.....	31

2.4.2 Services Provided by IP Mobility	33
2.4.3 Mobile IP Overview	33
2.5 IP Mobility Support for Multicast.....	35
2.6 The Mobile Multicast (MoM) Approach To Supporting Mobile IP Multicast.	36
2.6.1 MoM Description	36
2.7 MoM Simulation and Results.....	40
2.7.1.1 Simulation Model.....	40
2.7.1.2 Simulation Results	41
2.8 Summary	43
Chapter 3: Methodology.....	44
3.1 Introduction	44
3.2 Problem Overview	44
3.2.1 Problem Definition	45
3.2.2 Problem Statement	45
3.2.3 Problem Scope.....	45
3.2.3.1 Network Topology and Dimensions	46
3.2.3.2 : Number of Multicast Groups	47
3.2.3.3 Mobile Nodes	47
3.2.4 Method of Evaluation.....	47
3.2.5 Simulation Tool.....	48
3.3 Minimal Multicast Encapsulation	48
3.4 Operational Assumptions	52
3.4.1 Mobile IP Configuration Assumptions	52
3.4.1.1 Care-of address Assignment.....	52
3.4.1.2 Agent Advertisements.....	52
3.4.2 Network Components.....	53
3.4.2.1 Routers	53
3.4.2.2 Mobility Agents.....	53
3.4.2.3 Mobile Nodes	53
3.4.3 Network Links	54
3.4.3.1 Fixed Link	54
3.4.3.2 Mobile Links	54
3.4.4 Group Sources	54
3.4.5 Group Membership	55
3.4.6 Multicast Routing Algorithm Selection	55
3.4.7 Application Arrival Rate and Size	55

3.4.8 Background IP Traffic.....	55
3.4.9 Mobile Node Movement	55
3.5 Model Design and Operation	56
3.5.1 OPNET IP Node Modules.....	56
3.5.1.1 OPNET IP Process Modules	56
3.5.1.2 Changes to Existing OPNET IP Process Modules	57
3.5.2 New IP Node Models	60
3.5.2.1 Mobile Node Node Model.....	60
3.5.2.2 IP Mobility Support Modules.....	61
3.5.2.3 Mobile Node Link Processes.....	63
3.5.2.4 IP Mobility Agent Node Model	66
3.5.2.5 IP Mobility Support Modules.....	66
3.5.2.6 Link Processes.....	67
3.6 IP Mobility Multicast Support Mechanisms Tested.....	69
3.6.1 Multicast Join Mechanisms	69
3.6.2 Multicast Transmission	69
3.6.3 Mobile Join and Transmission Combinations	70
3.7 Experimental Factors	70
3.7.1 Mobile Group Size	70
3.8 Performance Metrics	70
3.8.1 Average Relative Path Length.....	71
3.8.2 Number of Lost Packets	71
3.9 Model Verification and Validation	71
3.9.1 Model Verification	71
3.9.2 Model Validation.....	72
3.10 Summary	74
Chapter 4: Results and Analysis.....	75
4.1 Introduction	75
4.2 Statistical Accuracy.....	75
4.3 Simulation Scenarios.....	76
4.3.1 Simulation Execution Length.....	76
4.4 Data Collection Methodologies.....	78
4.4.1 Path Efficiency	78
4.4.2 Packet Loss.....	79
4.4.3 Required Throughput	81

4.5 Analysis and Comparison of Performance Metrics	81
4.5.1 Analysis of Path Ratios	82
4.5.1.1 Hop Ratio	82
4.5.1.2 Distance Ratio	84
4.5.2 Analysis of Packet Loss	86
4.5.2.1 Degraded Link Change Proportion	86
4.5.2.2 Packet Loss Proportions	88
4.5.3 Analysis of Required Throughput	90
4.6 Summary	92
Chapter 5: Conclusions and Recommendations.....	93
5.1 Restatement of Research Goal	93
5.2 Conclusions	93
5.2.1 Results Synopsis.....	93
5.2.2 Recommendations	94
5.3 Significant Results of Research	95
5.4 Future Research.....	95
Bibliography.....	97
Vita	99

List of Figures

Figure 1: Simplified depiction of IP in IP tunneling	30
Figure 2: Framework Nations Network Topology	46
Figure 3: Original IP Header	51
Figure 4: Minimal Multicast Encapsulated Header	51
Figure 5: OPNET IP Router node model	58
Figure 6: OPNET Point-to-point Client and Server node models	58
Figure 7: Mobile Node Node Model	61
Figure 8: Link Packet Fields	64
Figure 9: Link ICI Fields	64
Figure 10: IP Mobility Agent Node Model	68
Figure 11: Path ratio collection methodology	80

List of Tables

Table 1: Permanently assigned low-level multicast groups	12
Table 2: TTL Scoping thresholds [Mau98].....	13
Table 3: Sample DVMRP routing table	24
Table 4: Sample DVMRP forwarding table	25
Table 5: PIM-DM Source-group Forwarding State Entry	28
Table 6: Mobile Join and Transmission Combinations.....	70
Table 7: Hop ratio results for Bi-directional tunneling with Home Tunneling.....	82
Table 8: Hop ratio results for Bi-directional Tunneling with Minimal Multicast Encapsulation	83
Table 9: Hop ratio results for Remote Subscription with Home Tunneling	83
Table 10: Hop ratio results for Remote Subscription with Minimal Multicast Encapsulation	83
Table 11: Maximum Observed Hop Ratios.....	83
Table 12: Distance Ratio Results for Bi-directional tunneling with Home Tunneling.....	84
Table 13: Distance Ratio Results for Bi-directional Tunneling with Minimal Multicast Encapsulation	85
Table 14: Distance Ratio Results for Remote Subscription with Home Tunneling.....	85
Table 15: Distance Ratio Results for Remote Subscription with Minimal Multicast Encapsulation	85
Table 16: Maximum Observed Distance Ratios	85
Table 17: Degraded Link Change proportion Results for Bi-directional tunneling with Home Tunneling.....	87
Table 18: Degraded Link Change proportion Results for Bi-directional Tunneling with Minimal Multicast Encapsulation	87
Table 19: Degraded Link Change proportion Results for Remote Subscription with Home Tunneling	87
Table 20: Degraded Link Change proportion Results for Remote Subscription with Minimal Multicast Encapsulation	88
Table 21: Packet loss proportion results for Bi-directional tunneling with Home Tunneling	88
Table 22: Packet loss proportion results for Bi-directional Tunneling with Minimal Multicast Encapsulation	89
Table 23: Packet loss proportion results for Remote Subscription with Home Tunneling	89
Table 24: Packet loss proportion results for Remote Subscription with Minimal Multicast Encapsulation	89
Table 25: Anomalous results for 120 Node Scenarios using Bi-directional Tunnel Combinations	90
Table 26: Maximum Observed Required Mobility Agent Throughput	91

Abstract

Given the expeditionary nature of current and future Air Force operations, the Air Force will continue to rely on ad hoc mobile networks to accomplish its operational objectives. Mobile multicast technology will provide two major benefits. First, it will allow for the efficient the use of "push" technology for dissemination of mission critical information, and second, it will provide for improved coordination and control of operations involving entities dispersed and moving throughout the battle space.

Allowing mobile hosts to connect to different links on an internetwork while keeping the same IP address is the challenge of IP mobility. In RFC 2002, "IP Mobility Support", the Internet Engineering Task Force (IETF) defines the protocols and mechanism to provide IP mobility. Additionally, the IETF defines two mechanisms that allow mobile nodes to transmit multicast packets, and two mechanisms that allow mobile nodes to receive multicast packets. The transmission mechanisms are direct transmission, and home tunneling. The reception mechanisms are bi-directional tunneling and remote subscription.

As proposed, the current IETF direct transmission mechanism requires the use of costly co-located care-of addresses, and does not provide for identification of senders' home IP addresses. This thesis proposes a novel modification to the IETF direct transmission mechanism. This modification, known as minimal multicast encapsulation, uses a modified form of minimal IP encapsulation to allow direct transmission while

utilizing less costly foreign agent care-of addresses. Minimal multicast encapsulation also provides positive identification of sender IP addresses.

In addition to demonstrating the viability of minimal multicast encapsulation, this research examines the performance of the four possible combinations of minimal multicast encapsulation or home tunneling with bi-directional tunneling or remote subscription. Comparisons are made in terms of path length, packet loss, and required mobility agent throughput. Results of this research indicate that, in terms of path efficiency, the combination of bi-directional tunneling with home tunneling suffers, on average, 3 times greater hop-based path length and between 4 and 5 times greater distance-based path length than the optimal combination of remote subscription with minimal multicast encapsulation. Also noted are extreme "worst cases" of 15 times hop length and 178 times distance-base path length. It is further demonstrated that bi-directional tunneling causes roughly 10 times more degraded links due to packet losses caused by delays inherent in the bi-directional tunneling mechanism. Finally, it is shown that bi-directional tunneling can increase maximum loading on mobility agents by up to 20 times over the loading experienced with remote subscription.

A COMPARATIVE ANALYSIS OF PROPOSED MOBILITY SUPPORT SCHEMES FOR IP MULTICAST

Chapter 1: Introduction

1.1 Background

Given the expeditionary nature of current and future Air Force operations, the Air Force will continue to rely on ad hoc mobile networks to accomplish its operational objectives. Mobile multicast technology will provide two major benefits. First, it will allow for efficient use of "push" technology for dissemination of mission critical information. Second, it will provide for improved coordination and control of operations involving entities widely dispersed and moving throughout the battle space.

Multicast communication involves communication from one-to-many or from many-to-many hosts. It allows a sender to transmit a data packet to a group of receivers by sending one packet to the group address. Multicast routers between the sender and the group of receivers replicate and route the packet as necessary to ensure delivery to all group members

Allowing mobile hosts to connect to different links on an internetwork while keeping the same IP address is the challenge of IP mobility. In RFC 2002, "IP Mobility Support", the Internet Engineering Task Force (IETF) defines the protocols and mechanism to provide IP mobility. Additionally, the IETF defines two mechanisms that allow mobile nodes to transmit multicast packets, and two mechanisms that allow mobile nodes to receive multicast packets. The transmission mechanisms are direct transmission, and

home tunneling. The reception mechanisms are bi-directional tunneling and remote subscription.

As defined, the current IETF direct transmission mechanism requires the use of costly co-located care-of addresses, and does not provide for identification of senders' home IP addresses. This thesis proposes *minimal multicast encapsulation*, a novel modification to the IETF direct transmission mechanism. Minimal multicast encapsulation uses a modified form of minimal IP encapsulation to allow direct transmission while utilizing less costly foreign agent care-of addresses. Minimal multicast encapsulation also provides positive identification of sender IP addresses.

1.2 Research Goal

This research has two goals:

- to introduce and determine the feasibility of a minimal multicast encapsulation
- to compare the performance of currently proposed IP mobility support mechanisms for IP multicast in terms of routing efficiency, packet loss and mobility agent loading.

1.3 Research Motivation

Research exists that analyzes IETF mobility support mechanisms for multicast operation. Harrison, et al [HaW97, ChW98, WiH98] introduced and analyzed a third mechanism for allowing mobile nodes to receive multicast traffic. This mechanism is called Mobile Multicast (MoM). MoM research primarily focuses on analyzing the performance of the MoM protocol, and while it does analytically compare the MoM

protocol to the currently proposed IETF mobile multicast support mechanisms, it abstracts away the supporting network topology and routing mechanisms that would, in reality, support such a system. The research does not consider multiple multicast sources, and does not allow the mobile hosts to be sources.

Harrison, et al compared the routing efficiency of the various mobility support mechanisms. While they admit that the combination of direct transmission with remote subscription provides optimal routing, they discount the use of the IETF direct transmission because it requires co-located care-of addresses for proper routing [HaW97, ChW98, WiH98]. This thesis proposes minimal multicast encapsulation as an alternative method of direct transmission that allows for optimal routing but does not require co-located care-of addresses.

This research effort analyzes the performance of minimal multicast encapsulation along with the currently proposed IETF mobile IP multicast support mechanisms. It expands upon previous research by examining the operation of the IETF support mechanisms under more realistic network conditions. Further, this research expands previous research by allowing more than one multicast source. Additionally, the mobile receivers of multicast are allowed to be multicast sources.

This research provides a performance analysis of four possible combinations of mobile multicast transmit and receive mechanisms. It determines which combination provides the best path efficiency and packet loss characteristics. It also examines the load placed on mobility agents while supporting visiting nodes.

1.4 Approach

This research OPNET Modeler to create realistic models of mobile nodes and IP mobility agents. These models are placed in a representative network environment. Supporting network infrastructure and protocols are modeled as well. The two IETF mobile multicast reception mechanisms are included in the models along with the IETF home tunneling transmit mechanism. The new transmit support mechanism, minimal multicast encapsulation, is also included in the model. Trial scenarios are executed using four possible combinations of transmit and receive mechanisms. Performance metrics of to path efficiency, packet loss and required mobility agent throughput are collected and analyzed, permitting conduct a comparative analysis of the four combinations.

1.5 Overview of Results

This thesis documents the inherent inefficiencies in the bi-directional tunneling, and home tunneling multicast mobility support mechanisms. The combination of bi-directional tunneling and home tunneling proved to be the least efficient in terms of path length; the average path length (in hops) was, on average, three times longer than the optimal path, and the average physical path length was from thee to four times longer than optimal. The combination of remote subscription and minimal multicast encapsulation proved to be the most efficient in terms of hops and distance; the average ratio for both of these metrics was one. Bi-directional tunneling also shown to suffer greater packet loss than remote subscription because of inherent delays associated with updating the home agent. Bi-directional tunneling resulted in up to 20 times higher loading on mobility agents than did remote subscription. Finally, this research

demonstrates that while causing a slight increase network loading due to encapsulation, minimal multicast encapsulation is indeed a viable direct transmission alternative that provides optimal multicast routing.

1.6 Summary

This chapter introduced the background and goals for this research. It went on to explain the motivations for this research and presented an overview of the results of this study. The remainder of this thesis is organized as follows: Chapter 2 presents a review of current literature and research in the areas of IP multicast and IP mobility. Chapter 3 provides an overview of the methodology used to construct models needed to perform the analysis of the mobile multicast support mechanisms. Additionally in Chapter 3, Section 3.3 provides a detailed description of minimal multicast encapsulation, the novel direct transmission mechanism introduced in this work. Chapter 4 presents the data collection techniques used in this thesis along with an analysis of the result obtained through network simulation. Chapter 5 concludes the thesis with a discussion of conclusions and recommendations for areas of future research in mobility support for IP multicast.

Chapter 2: Literature Review

2.1 Introduction

This section presents a review of the current literature and research in the area of Internet Protocol (IP) mobility support for multicast. It begins with an introduction to the current state of IP multicast. The IP multicast group management (IGMP) protocol is discussed along with several multicast routing protocols and implementations. Next IP in IP tunneling is introduced as a prelude to IP mobility support. IP mobility support is then presented with particular emphasis on the IETF-proposed methodologies for mobility support of multicast. Finally an alternative to the IETF methodologies is presented along with a discussion of initial performance analyses conducted comparing the alternative methodology to the proposed IETF methodologies.

2.2 IP Multicast

2.2.1 Multicast Introduction

IP multicast communications and associated applications are on the verge of becoming the next great internetworking technology. Multicast communication involves one or many senders sending data to many receivers who wish to receive the particular data. Broadcast communication, in contrast, involves one sender sending data to all receivers on a network. Unicast, the prevalent internetworking communication technology today, is the sending of data from one sender to one receiver.

The one-to-many or many-to-many effects of multicast can be achieved through either broadcast or unicast but at greater costs. To achieve these effects with unicast, every sender must be supplied with and maintain a list of the receivers' addresses that wish to receive (or subscribe to) its transmission. The sender must then send out an individual stream of data to each subscriber. This means that for n subscribers, the unicast transmitter would have to send n copies of the message. Using broadcast, the sender would not have to keep track of its recipients, but the message would be sent to every subnet and every node on the network, regardless of whether there were any hosts on a given subnet that wished to receive the broadcast message. With either solution, unicast or broadcast, unnecessary replication or transmission of data occurs.

The goal of multicast technology is to minimize duplication and transmission of data so that data travels only along network paths used to get to desired destinations. Multicast provides the concept of multicast group addresses. This means that multicast senders (sources) do not need to know the individual address of each receiver (subscribers). Sources just need to know the multicast group address to which users are subscribed, or users need to know to which group address a source sends so that they can subscribe and listen to that particular multicast group address.

Multicast applications make use of multicast communication to carry out required tasks. Multicast applications are capable of true "push" communication in which a source sends (pushes) data to all subscribed hosts as it becomes available on the source without having to be continuously queried by individual subscribers for information.

Multicast provides scalability for these types of push-based applications. Scalability is the ability to handle an increased number of users without incurring detrimental effects on the network or servers [Mil99].

An example of how this type of application does not scale well with unicast protocols is the commercial, web-based PointCast application [Mau98]. PointCast was a fairly recent (1996) attempt to emulate push technology through what amounted to unicast-based, automated pull technology. In the PointCast system, subscriber hosts automatically would periodically request transmissions of data updates from source hosts. The source hosts would then send out data streams to requesting subscribers individually as requests came in. This became a problem on many corporate LANs because many subscriber hosts were requesting the same information from the same source hosts, which flooded the LANs with duplicate traffic [Mil99].

Many other applications are also better suited for multicast than unicast. These applications include large-scale software distribution and updating, teleconferencing, "broadcast" entertainment or educational applications, collaborative group applications such as distributed electronic white boards, and distributed interactive simulations. Military command and control networks could also benefit from multicast technology by enabling up to the minute intelligence information to be pushed from many diverse sources to all units requiring the information. Multicast technology would also improve the synchronized coordination and control of geographically dispersed units.

Since IPv4 is currently the most ubiquitous network protocol, this research will deal with IPv4 based multicast. In the remainder of section 2.2, IP multicast is introduced along with a discussion of the IPv4 multicast addressing scheme, administrative scoping,

Internet Group Management Protocol, various multicast forwarding algorithms, and the current implementations of these algorithms.

2.2.2 Host Extensions for IP Multicasting

The Internet Engineering Task Force (IETF) Request for Comment (RFC) 1112, "Host Extensions for IP Multicasting," specifies the required extensions for "a host implementation of the Internet Protocol (IP) to support multicasting." The IETF further defines RFC 1112 as "... the recommended standard for IP multicasting in the Internet" [Dee89].

RFC 1112 describes IP multicasting as the transmission of IP datagrams to a "host group" in which a set of zero or more hosts is designated by a single IP destination address. These multicast datagrams are delivered to all members of (or subscribers to) the destination host group with the same "best efforts" reliability as regular unicast IP datagrams. "Best-efforts" reliability means the IP datagram is not guaranteed to reach all destination group members and that the datagrams may arrive out of order [Dee89].

Membership in host groups is dynamic. This means that at any given time, a host group may have zero or more members. The members of, and sources to a host group may be located anywhere on the internetwork. Members may belong to more than one host group at a time. A source host does not need to be a member of a host group to send datagrams to that group [Dee89].

Host groups may be either permanent or transient. Permanent groups are administratively assigned well-known IP address. The term permanent refers to the fact that the address is permanently assigned and not the number of members hosts

subscribed. Permanent groups continue to exist even if no hosts are subscribed to them. Transient groups, on the other hand, are assigned addresses not reserved for permanent groups. The addresses are assigned on an as needed basis and cease to exist when no hosts are subscribed to them [Dee89].

Multicast routers are routers that forward multicast datagrams for a particular group to local hosts that are host group members, and to other areas of the network that contain host group members. Multicast routers and hosts must also employ the Internet Group Management Protocol (IGMP) in the IP layer. This allows hosts supporting IGMP to notify multicast routers of their membership in host groups. The multicast router may then send its membership information to other multicast routers with members of, or sources for, that host group in their subnets. In this way, multicast trees can be created to ensure connectivity among all host group members and all sources wishing to send to that group. These senders need not necessarily be group members.

2.2.3 Addressing

Host group addresses comprise the set of IP addresses known as class D addresses. Class D addresses have “1110” as their high order bits and thus range from 224.0.0.0 to 239.255.255.255. The Internet Assigned Numbers Authority (IANA) assigns permanent groups. The IANA publishes the list in its “Assigned Numbers” registry. The addresses from 224.0.0.0 to 240.0.255 are reserved for exchange of routing information and other low-level topology discovery or maintenance protocols, such as gateway discovery and group membership reporting. Multicast routers should not forward datagrams addressed

to these groups. Table 1 lists the 19 low-level group address groups that have thus far been designated by the IANA.

Currently there are over 4000 permanent multicast group addresses assigned to organizations and corporate groups in the address ranging from 224.0.1.0 to 224.0.22.255. For a complete listing of all assigned multicast group addresses, the reader is referred to the IANA web site <http://www.iana.org>.

Another important address group is composed of the administratively scoped multicast group addresses which range from 239.0.0.0 to 239.255.255.255, which is subdivided into various levels of scoping which range from site-local scope to organizational-local scope. The purpose of administratively scoped addresses is to provide organizations with addresses that may be assigned internally at an organizational or site level for multicast groups that will not extend beyond the organizational or site boundaries. Multicast routers at the borders of such boundaries should not forward multicast packets addressed to groups of lesser or equal administrative scope than the border upon which they reside. The boundary regions defined by the boundary routers must be convex boundaries. This means that no path between two internal non-border routers within the same administrative region can pass outside the administrative region [Mau98].

Table 1: Permanently assigned low-level multicast groups

224.0.0.0	Base Address (reserved and guaranteed to be unassigned)
224.0.0.1	All systems on this subnet
224.0.0.2	All routers on this subnet
224.0.0.3	Unassigned
224.0.0.4	DVMRP Routers
224.0.0.5	OSPFIGP OSPFIGP All Routers
224.0.0.6	OSPFIGP OSPFIGP Designated Routers
224.0.0.7	ST Routers
224.0.0.8	ST Hosts
224.0.0.9	RIP 2 Routers
224.0.0.10	IGRP Routers
224.0.0.11	Mobile-Agents
224.0.0.12	DHCP Server / Relay Agent
224.0.0.13	All PIM routers
224.0.0.14	RSVP-Encapsulation
224.0.0.15	All CBT routers
224.0.0.16	Designated-sbm
224.0.0.17	All-sbms
224.0.0.18	VRRP

Administratively scoped addresses alleviate the problems caused by the Time To Live (TTL) scoping scheme. The IP TTL field contains a threshold value that determines the scope of a multicast datagram. TTL scoping thresholds are shown in Table 2. The problem with relying on TTL thresholds was that the TTL sometimes would not be set high enough to reach an interested outsider, or it might not be set high enough to cross a threshold boundary. For example, if the TTL threshold was set to the 16 so that it would get outside the local site, but for some reason made more than one hop inside the local site, it would not get past the local site boundary router. At the other extreme, if network

administrators were not careful enough and calculated an overly large TTL value, private multicast information could leak out past the site or organizational boundaries [Mau98].

Table 2: TTL Scoping thresholds [Mau98].

TTL	Scope Threshold
0	Restricted to same host
1	Restricted to same subnet
15	Restricted to same site
63	Restricted to same region
127	Worldwide; limited bandwidth
191	Worldwide
255	Unrestricted Scope

2.2.4 IGMP

For a host to receive multicast traffic destined to a given host group, it must inform its local multicast router that it wishes to join the given group. Also according to RCF-1112, at least one multicast router per LAN must periodically query the hosts on the LAN to see if they are still members of a multicast group. IGMP provides the mechanism for local hosts to inform their designated multicast router that they want to join or are still members of a multicast group. Presently two versions of IGMP have been released. IGMPv1 was defined in the appendix of RFC 1112. IGMPv2 was published as RFC 2236, a proposed standard, updating the standard set forth in RFC 1112. Currently a newer version, IGMPv3, exists only as an Internet Draft.

Currently, IGMPv2 is the most wide spread version. It shares many features with v1 and introduces some enhancements. The features of v1 are presented followed by a discussion of v2 enhancements.

As stated above, IGMP provides a mechanism for a multicast router to query the hosts on the directly attached network to see if they are members of any multicast group.

This is accomplished by sending out a query message to the “All hosts” multicast group 224.0.0.1 with a TTL of 1. This ensures that only hosts on the directly attached network will respond to the router’s query.

Upon receipt of the IGMP query message, a host transmits a membership report to the group addresses of each group for which he is a member. This response also has a TTL value of 1 to prevent it from going past the router. It is important to note that in both v1 and v2, these membership reports are addressed to the multicast group addresses of which the host is a member and not to the “All routers” address. This means that the router must listen for all multicast IP addresses. Since the router must already do this to route multicast messages, there is no additional burden for the router.

To keep from flooding the router with membership requests immediately after sending a query, IGMP provides for a random wait time before each host answers a query. During this wait time, the hosts listen to the multicast group to which they would like to subscribe or already belong. If they hear another system answer the query to one of those groups, they do not respond. This limits the number of query response messages that are sent. The member hosts are still ensured of having multicast messages forwarded to them because the router need only know that there are one or more host members of a group on its subnet in order to forward packets for that group. When hosts wish to leave a group under IGMPv1, they simply stop responding to the periodic router queries. If after several queries, the router receives no responses back from hosts for subscribing to a particular group, it assumes that no more hosts are subscribing to the group from its subnet and stops forwarding multicast packets for that group to its subnet.

To avoid the latency involved in waiting to be queried by the router, hosts send a report message as soon as they want to join a group. This reduces the join latency experienced by the member host. The join latency is defined as the time it takes from when the first membership report is issued by a host to the time when that host starts receiving traffic destined to the group. This latency can be as low as zero if the group is already active on the local network. If the group is not active on the local network, the join latency is the time that it takes the multicast router to connect to the group's multicast tree (assuming that it is active somewhere on the internet).

Only one multicast router can be active on a subnet at a given time. IGMPv1 requires that the routing protocol elect a querier, but IGMPv2 provides an election algorithm for choosing the querier. The election algorithm is simple: all multicast routers start up as queriers for their networks, and if they hear another router querying with a lower IP address, they must become non-queriers [Fen97].

Another primary feature that IGMPv2 adds to v1 is the "leave group" message. If a host desires to leave a group and it was the last host to transmit a group membership report, it should transmit a "leave group" message to the 224.0.0.2 "All routers" address. If it was not the last host to transmit a group membership report, then it assumes there is another group member on the network and the leaving host does not need to transmit anything when it leaves the group. If a host does not have enough memory to remember whether it was the last host to transmit a membership report, it may transmit a leave group message whenever it wishes to leave a group [Fen97].

Upon receipt of the "leave group" message, the multicast router sends a group specific "last member count" query to the group of which the host that just left was a

member. The router waits an adjustable (default 1 second) time interval before it decides that no hosts are subscribed to the group and stops forwarding traffic to the group. This adjustability allows network administrator the ability to tailor the leave latency time [Fen97].

The leave latency is the time that it takes from when the last host on the network left the group by transmitting a “leave group” message to when the router stops forwarding messages to the group on the LAN. With IGMPV1 the latency could be quite long because the router would have to wait for several unanswered membership queries (usually 3) before deciding that no host were left for a given group. This time could be on the order of minutes. By having the adjustable threshold for time after the “Last member count” query, the leave latency can be made much lower. Lowering the leave latency can be very beneficial when a group is receiving a heavy load of traffic, because the sooner the router can remove itself from the group’s routing tree, and relieve itself of the load, the better [Mau98].

2.2.5 Forwarding Algorithms

The IGMP is merely a local protocol used to ensure that multicast routers know that their subnets contain group members. This ensures that once the router receives packets destined for a given group, it can forward them to group members on its subnet. Different protocols are required to ensure that multicast packets are routed throughout the internetwork so that all subnets with member hosts receive packets destined for their designated multicast groups. There are two basic types of multicast algorithms: dense-mode and sparse-mode. Dense mode algorithms involve a form of flooding at some

point, while sparse mode algorithms use selective techniques to create and maintain multicast trees [GoN99]. The terms dense and sparse refer to the distribution of subscribers throughout the network. If subscribers can be found on most subnets in an internetwork, then they are considered to be densely distributed. If, on the other hand, subscribers are located at a relatively small proportion of subnets, they are considered to be sparsely distributed. Note, however, that the terms dense and sparse do not imply anything about the total number of subscribers present in an internetwork.

Another way to categorize multicast forwarding algorithms is by whether they are source-based or shared-tree. The source-based algorithms construct distinct delivery trees based on source location while shared-tree approaches generally use only one delivery tree independent of source [Mau98].

Of the currently implemented multicast forwarding protocols, some are source-based and some are shared-tree based. Of these protocols, dense mode protocols can be either source based or shared-tree based, but the sparse-mode tend to only be shared-tree because of the poor scaling capacity of source-based protocols. Source-based protocols scale less readily than shared-tree protocols because source-based protocols must maintain path information for each source-group pair while shared-tree protocols must only maintain an entry for each group [Mau98].

To understand the currently implemented multicast forwarding protocols, it is useful to look at the evolution of multicast algorithms from the simplest broadcast algorithms such as flooding to the two more current algorithms: reverse path multicasting and core-based trees.

2.2.5.1 Flooding

Flooding is the simplest multicast routing algorithm. Flooding involves sending packets to all routers on a network. In a flooding algorithm, a router receives a multicast packet and, if it hasn't recently seen the packet, forwards it along all outgoing paths except the one it came in on. If it has recently seen the packet it discards it. Flooding guarantees that all routers will eventually get the packet whether they have subscribers on their subnets or not.

There are several obvious drawbacks to the flooding approach. First, it does not scale well in wide area networks. Second, it sends packets to router that have no hosts interested in receiving them. Third, it is an inefficient use of router memory, because a router must maintain information about recently seen packets in a table. At high data rates, this could mean numerous table entries [GoN99].

2.2.5.2 Spanning Trees

Tree construction is an efficient method for delivering multicast packets. A spanning tree is a structure in which a single path connects any two routers on a network. Spanning tree algorithms are very useful in removing loops from the network topology. Spanning trees are set up by ensuring that each router has only one port to reach a designated subnet on the internet. Then at each subnet, each designated root node is assigned and is configured ensure that only one path exist between itself an all other routers on its subnet. Spanning tree algorithms are powerful in removing loops and fairly easy to implement, but tend to cause centralization of traffic because all traffic on the network only travels across the spanning tree's edges [GoN99].

The algorithms that follow all create some sort of spanning tree (multiple or singular). The goal of these algorithms is to keep every router from having to compute and maintain the entire spanning tree. It is also desirable to limit the spanning tree in such a way that it spans only nodes that have member hosts on directly connected subnets and those nodes interconnecting them.

2.2.5.3 Reverse Path Forwarding

Reverse path Forwarding (RPF) slightly improves efficiency over simple flooding, by eliminating the need to keep track of recently received packets in order to eliminate loops. RPF does this by checking whether a received packet came in on a parent interface that is along the shortest path back to the source. If it was, the packet is broadcast on all child interfaces (all interfaces except the one upon which the packet arrived). If the parent interface was not on the shortest path back to the source, it discards the packet. RPF suffers from the fact that a packet will be forwarded to a given node, and thus its subnets, as many times as it has parent nodes since all of its parent nodes will eventually receive the packet along what they consider to be their shortest path back to the source [DeC90].

2.2.5.4 Reverse Path Broadcast

Reverse path broadcast (RPB) is actually a broadcast algorithm, and is the forerunner of reverse path multicast. RPB is similar to RPF except that it ensures that all packets are sent out only on child interfaces that are truly “down stream” from the source. This means that a packet is sent to a child if the child interface is also on the child’s shortest path back to the source. If a router is utilizing a link state routing algorithm, it can easily

determine if it is on a child node's shortest path back to the source. If it is utilizing a distance vector routing algorithm, child routers can advertise their previous hop information to a parent router in routing messages.

RPB is effective in ensuring that all nodes receive packets only once. This algorithm is efficient as it guarantees the shortest delivery path from the source to each recipient. It is still a broadcast algorithm and does not take into account group membership during forwarding. All routers and their attached subnets get copies of all multicast packets regardless of whether there are any member hosts on the subnet.

2.2.5.5 Truncated Reverse Path Broadcast

Truncated reverse path broadcasting (TRPB) is a further improvement on RPB in that routers do not forward packets to their connected subnets if the subnets do not contain any member hosts for a given group. TRPB routers would use IGMP information in deciding whether to truncate the path (i.e., to not send packets to a connected subnet). TRPB thus reduces the load on uninterested subnets, but still does not make use of group membership information in deciding whether to forward packets to downstream routers. It simply forwards to all valid downstream router as in RPB [Mau98].

2.2.5.6 Reverse Path Multicasting

Reverse path multicasting (RPM) carries TRPB one step further. The algorithm uses IGMP host membership information to "prune" nodes off the multicast distribution tree. It does this by sending the first packet destined to a group to all nodes as in TRPB. Then, however, leaf nodes that do not have any member hosts connected send prune messages to their parent routers. Upon prune message receipt, the parent router stops forwarding

traffic for the designated group to the child router that sent the prune message. If all of a parent router's child routers send prune messages for a particular group, then the parent router sends a prune message to its parent router thereby possibly pruning entire sub-trees. If a child later discovers that it has member hosts for the pruned group, it simply issues a graft message to its parent. This causes the parent to resume forwarding packets, if the parent itself has not issued a prune message. If it has issued a prune message, the parent then sends a graft message to its parent. This process continues until the sub-tree has been grafted back to the multicast distribution tree.

To keep from having to store prune state information indefinitely for a group (if no one is sending to the group), the prune information is aged and periodically deleted from router memory. If the group still has active sources, all links whose prune information has expired will again receive at least one packet for the given group. The routers will then have to reinitiate the pruning process if they still do not have directly connected member hosts or downstream children with the same [Mau98].

2.2.5.7 Core-Based Trees

As previously stated, the primary problem with source-based tree algorithms such as RPM lack of scalability for as the number of sources and/or the number of groups increase. This is because router must keep routing and forwarding information for each source-group pair. Also, the periodic broadcast of multicast traffic to all nodes is undesirable especially in a sparse host member environment.

To solve the problem of operating in a sparse environment and limit the amount of multicast routing and forwarding information stored, Core Based Trees (CBT) were

developed. With CBT, one core router (CR) is chosen for a given group within a CBT domain. When a CBT-aware router discovers that it has a host member wishing to join a group, the router sends a Join Request message in the group CR's direction. If an intermediate router is not on the group's core tree, it forwards the Join Request in the CR's direction either until the request reaches the CR or an on-tree router.

The intermediate router also stores state information regarding the Join Request path back to the requestor. This is done to allow for the addition of two interfaces to the group forwarding cache. The two interfaces added are for the request the acknowledgement received. If an intermediate router is already connected to the distribution tree, it will send a join acknowledge back to the requestor and add the interface upon which the request came to its forwarding cache for that particular group. The CR behaves in the same manner if the request message reaches it.

Finally, at the requestor, the interface upon which the acknowledgement was received is also added to its forwarding cache. In this way, when a router receives a multicast transmission for a group on any of its on-tree-interfaces, it simply needs to forward it on all of its other on-tree interfaces to ensure that the message reaches all nodes in the tree.

While CBT routers do not need to care whether a packet came in on the shortest path back to the source, they must know which interface leads back to the CR. This interface is considered their upstream interface and routers are responsible for making sure that it is operational by periodically sending echo requests to their upstream router. If an echo request does not get answered, the router assumes that the upstream link is non-operational. It then must issue a flush message to all of its downstream routers who must then do the same for their downstream routers. Upon issuing or forwarding a flush

message, a router must delete all group state information for that particular group and attempt to reestablish contact with the core tree as described above [Bal97].

In CBTs, off-tree (non-member host) sources must unicast an encapsulated (IP over IP) multicast message directly to the CR, which will then forward the unencapsulated multicast message on all of its on-tree interfaces for the group.

For CBTs to function, CBT aware routers must be able to determine which router in a CBT domain is the CR. This is accomplished by a bootstrap protocol. In the bootstrap protocol, network administrators determine which of their routers will be Core Candidates (CCs). Once a router is designated as a CC, it notifies the previously agreed upon Bootstrap Router (BR) for the CBT domain. The CC includes in its notification, the groups for which the CC is willing to act as a CR. Once the BR has a list of all CCs, it multicasts this list along with the groups they are willing to represent to all CBT routers via the "All CBT routers" multicast address. After the routers receive the list of CCs, they apply a hashing function based on the group address to come up with a list index value for that group's CR. The hashing function is designed so that a minimal number of consecutive group addresses are mapped to a given CR. All routers have the same hash function and CC list so all routers will come up with the same CR from the list of CCs. If the CR goes down, the BR is notified and re-collects a list of CCs. The new list is again sent out, and the hashing function is reapplied to find the new CR [Bal97].

2.2.6 Current Implementations

Current multicast implementations utilize variations of the previously described algorithms. This section describes four implementations. The first two are protocol

dependent. To provide multicast forwarding, routers are required implement specific multicast routing protocols. These specific multicast routing algorithms are implemented in addition to any unicast routing protocol that a router may utilize. The last two are protocol independent. They do not provide or rely on any specific multicast routing protocol. To make multicast forwarding decisions, they make use of whatever unicast routing protocol a given router implements.

2.2.6.1 Distance Vector Multicast Routing Protocol

The Distance Vector Multicast Routing Protocol (DVMRP) uses distance vector routing information to calculate the shortest reverse path for each source-group pair. A separate routing table must also be maintained. As Table 3 illustrates, a DVMRP routing table contains the following information: the source prefix, the subnet mask for the source prefix, the from gateway, the distance metric value, the link status, and the TTL. Note that the TTL field is used for table management and has nothing to do with packet TTL fields or scoping. This field simply indicates the number of seconds remaining before the table entry is invalid.

Table 3: Sample DVMRP routing table

Source prefix	Subnet mask	From gateway	Metric	Status	TTL
128.1.0.0	255.255.0.0	128.7.5.2	3	Up	200
128.2.0.0	255.255.0.0	128.7.5.2	5	Up	150
128.3.0.0	255.255.0.0	128.6.3.1	2	Up	150
128.4.0.0	255.255.0.0	128.6.3.1	4	Up	200

The DVMRP routing table does not contain any information about group membership and therefore, a forwarding table must be constructed. The forwarding table, shown in

Table 4, contains the source prefix, multicast group, parent interface (InIntf), and the child interfaces (OutIntf(s)). The InIntf field contains the interface number for the shortest path interface, and a “Pr” next to the entry indicates that a prune message has been sent to the router’s parent for that group. The OutIntf field contains the downstream child interfaces for that group, and “p” next to this entry indicates that a prune message has been received on that interface.

Table 4: Sample DVMRP forwarding table

Source Prefix	Multicast Group	InIntf	OutIntf(s)
128.1.0.0	224.1.1.1	1Pr	2p3p
	224.2.2.2	1	2p3
	224.3.3.3	1	2
128.2.0.0	224.1.1.1	2	2p3

When a multicast packet comes in from a source for the first time, the multicast routing table is first checked to see if the packet indeed came in from the correct shortest path gateway. If it did, then a forwarding table entry is created for the appropriate source-group pair with proper local port values for the InIntf and OutIntf(s). When subsequent packets come in for that source-group pair, the forwarding table is used to quickly determine whether to forward the packet. Also, the interfaces on which to forward them are determined according to the RPM algorithm [Mau98].

DVMRP has several limitations. Since it uses a distance vector protocol, the time to converge can be long. A second limitation is that the useable local network diameter is limited to approximately 15 hops [Mau98]. Thirdly, all routers must maintain source-group state information even when they are not on the tree. This causes additional problems, as the source state information does not scale well with increased number of

sources. As a final limitation, the multicast traffic is periodically broadcast across the entire network. This last problem is common to all RPM-based algorithms.

2.2.6.2 Multicast OSPF

Multicast OSPF (MOSPF) is an extension to the Open Shortest Path First unicast algorithm. In MOSPF, every multicast router in an OSPF region contains the state of every other router in the OSPF region. This is accomplished through link state announcements (LSAs) as in OSPF. In MOSPF, however, these LSAs contain group membership information for each router. In this way, all routers have, in a sense, a shared database of state information for the entire region. With MOSPF shortest path routes for each source-group pair are calculated using Dijkstra's algorithm. These paths, however, are not calculated until the first time a packet from a particular source-group pair is received. When this occurs, the proper incoming and outgoing interfaces are calculated for the given shortest path tree and entered into a forwarding cache. From this point on, the forwarding cache is used to correctly forward packets. The forwarding cache does not age, but is updated as required when new LSAs come in indicating either a change in network status or a change in group membership or source topology.

The primary weakness of MOSPF comes from the fact that Dijkstra's algorithm must be run on each node in the tree whenever a new source-group pair is added. This would be collectively quite computationally intensive in a highly dynamic environment.

2.2.6.3 Protocol-Independent Multicast

Protocol Independent Multicast (PIM) is actually comprised of two quite different algorithms: one for dense mode multicast (PIM-DM) and one for sparse mode multicast

(PIM-SM). They are similarly named, however, because they share common control messages.

2.2.6.3.1 PIM-DM

PIM-DM is an enhanced RPF algorithm that does not have its own routing protocol, and relies on the information in the unicast routing table to orient itself with respect to sources. When a packet arrives on an interface, the PIM-DM process utilizes the router's unicast routing protocol to perform a reverse-path routing lookup to the source. If the interface upon which the packet arrived is the same interface that the unicast routing protocol would use to send packets to the source, the interface is considered to be the correct incoming interface for that particular source. If the packet arrived on the correct incoming interface, it is then forwarded on all other interfaces that have PIM-DM routers or subscribed hosts attached. If the interface was not the correct interface, the packet is discarded and a PIM-Prune message is sent on the receiving interface. Since PIM-DM does not have a separate routing protocol, as do DVMRP and MOSPF, PIM-DM has no prior knowledge about whether an outgoing interface is connected to a downstream child. Because of this, packets are initially sent out on all interfaces that have PIM-DM routers or subscribers attached. DVMRP and MOSPF both use their built-in multicast routing protocols to avoid this situation. The designers of PIM-DM have chosen simplicity of operation over the overhead generated by initially forwarding packets to non-child routers [Dee99]. Like DVMRP, PIM-DM maintains a source-group forwarding state for every multicast source-group pair. Unlike DVMRP, however, the forwarding state entry is not created until packets are received from a particular source for a given group. Table

5 illustrates a source-group forwarding state entry. PIM-DM routers maintain source group entries as long as they continue to receive packets from the source. If the router does not receive a packet from the source for 210 seconds, it deletes the entire source-group entry [Dee99].

Table 5: PIM-DM Source-group Forwarding State Entry

Field	Value(s)
Source Address	192.34.5.3
Group Address	224.100.30.2
Outgoing Interface List	1,3,5

When a PIM-DM router receives a multicast packet, and no longer has subscribers or other PIM-DM routers attached to any of its outgoing interfaces, or if the outgoing interfaces list is empty, the router sends a PIM-Prune message to its upstream neighbor. If the interface upon which the upstream neighbor receives the PIM-Prune message is a point-to-point interface, the upstream neighbor immediately prunes the interface for the particular source-group combination. If, however, the interface upon which the upstream neighbor received that PIM-Prune message is connected to a multi-access LAN, the upstream router delays pruning the interface for 3 seconds. During this delay, any other PIM-DM router attached to the LAN that still wishes to receive packets from the upstream router must multicast a PIM-Join message to the “All PIM routers” multicast group. Upon receiving the join message, the upstream router, as well as router requesting the prune, cancel any pruning actions. This process is known as “prune override” [Dee99].

When an interface is pruned, it is removed from the forwarding state entry's outgoing interface list, and a 207 second prune state timer is started [Dee99]. When the prune state timer expires, the pruned interface is returned to the source-group entry's outgoing interface list. Down stream routers again receive multicast packets for the given source-group pair even if they still do not need to receive packets for the group. Upon receiving packets from a previously pruned link, routers that still do not need to receive packets must again send prune messages upstream.

2.2.6.3.2 PIM-SM

PIM-SM is based on the CBT algorithm with enhancements that allow shortest path trees to be constructed for each source. In PIM-SM the term core router is replaced by rendezvous point (RP) but serves the same purpose and is elected in the same manner. CBT Join messages are replaced by PIM Join messages and both member hosts and RPs are free to connect to the source's shortest path tree sending PIM-Join requests to the shortest path tree instead of the RP. This may occur when the source's shortest path tree has better bandwidth than the portion of the core tree that the RP or member host is using.

2.3 IP in IP Tunneling

To understand some the concepts presented in the discussion of mobile IP and mobile multicast, it is first necessary to understand the concept of IP in IP tunneling. Alternatively referred to as IP in IP encapsulation, several RFCs deal with the subject, the most recent being the proposed standard, RFC 2003, "IP Encapsulation Within IP" [Per96b]. IP in IP encapsulation was proposed to allow a specific route or tunnel for IP traffic that differs from the standard destination-based route. Of specific interest is the

use of IP in IP tunneling for delivery of datagrams to mobile hosts located on distant networks.

The IP in IP tunnel is defined by its two endpoints. At one end, a router or host encapsulates the original IP datagram within a new datagram containing the far tunnel endpoint address as the destination address. The far endpoint of a tunnel is simply the router that decapsulates the original IP datagram and delivers the datagram to its final destination using standard IP routing methods. The encapsulation causes the original datagram to follow a path that it would originally not follow as illustrated in Figure 1.

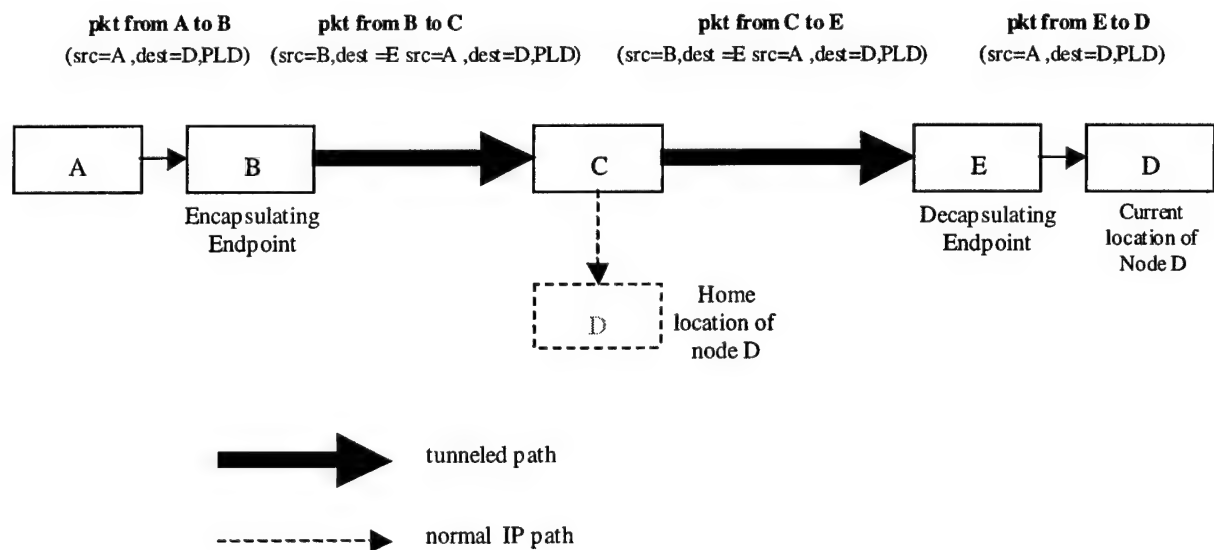


Figure 1: Simplified depiction of IP in IP tunneling

2.4 Mobile IP

As its name implies, Mobile IP is the network (IP-layer) solution to the problem of node mobility. First, however, we must define mobility. Mobility is the “ability of a node to change its point of attachment from one link to another while maintaining all existing communications and using the same IP address at its new link” [Per96a].

2.4.1 Mobile IP Definitions

RFC-2002, “IP Mobility Support”, defines the mechanisms and protocols by which nodes may implement mobility, and thus be considered mobile nodes. Prior to proceeding with a general description of the Mobile IP mechanism, a few definitions from RFC 2002 [Per96a] are required:

Mobile Node

A host or router that changes its point of attachment from one network or subnet to another. A mobile node may change its location without changing its IP address; it may continue to communicate with other Internet nodes at any location using its (constant) IP address, assuming link-layer connectivity to a point of attachment is available.

Home Agent

A router on a mobile node’s home network that tunnels datagrams for delivery to the mobile node when it is away from home, and maintains current location information for the mobile node.

Foreign Agent

A router on a mobile node’s visited network which provides routing services to the mobile node while registered. The foreign agent detunnels and delivers datagrams to the mobile node that were tunneled by the mobile node’s home agent. For datagrams sent by a mobile node, the foreign agent may serve as a default router for registered mobile nodes.

Correspondent Node

A peer with which a mobile node is communicating. A correspondent node may be either mobile or stationary.

Foreign Network

Any network other than the mobile node's Home Network.

Home Address

An IP address that is assigned for an extended period of time to a mobile node. It remains unchanged regardless of where the node is attached to the Internet.

Home Network

A network, possibly virtual, having a network prefix matching that of a mobile node's home address. Note that standard IP routing mechanisms will deliver datagrams destined to a mobile node's Home Address to the mobile node's Home Network.

Link

A facility or medium over which nodes can communicate at the link layer. A link underlies the network layer.

Link-Layer Address

The address used to identify an endpoint of some communication over a physical link. Typically, the Link-Layer address is an interface's Media Access Control (MAC) address.

Mobility Agent

Either a home agent or a foreign agent.

Mobility Binding

The association of a home address with a care-of address, along with the remaining lifetime of that association.

Node

A host or a router.

Tunnel

The path followed by a datagram while it is encapsulated. The model is that, while it is encapsulated, a datagram is routed to a knowledgeable decapsulating agent, which decapsulates the datagram and then correctly delivers it to its ultimate destination.

Visited Network

A network other than a mobile node's Home Network, to which the mobile node is currently connected.

2.4.2 Services Provided by IP Mobility

From the above-listed definitions, it is evident that IP mobility provides mobility agents to service the mobile nodes. Mobile nodes find mobility agents through an agent discovery service that allows for the mobile agent to solicit for and discover mobility nodes on the network to which they are currently attached. Using this service, mobility agents periodically announce their presence on a network, and newly arriving mobile nodes can request mobility agent services [Per96a].

Another service provided by IP mobility is the registration service. This service allows a mobile node to register with its home agent to inform the home agent of its new care-of address. A mobile node can register directly with its home agent or through a foreign agent [Per96a].

2.4.3 Mobile IP Overview

The following is an overview of the mobile IP protocol presented in RFC 2002 [Per96a]. Mobility agents (home or foreign) periodically advertise their presence on their respective networks via Agent Advertisement messages. If a mobile node has recently arrived, it can also solicit mobility agent services via an Agent Solicitation message, triggers mobility agents on the link to advertise their presence.

Once a mobile node receives an Agent Advertisement message, it can determine whether it is at its home network or at a foreign network. If the mobile node is located at its home network, it communicates as a non-mobile node would without using mobility services. One exception is for periodically receiving its home agent's Agent Advertisement messages to verify that it is indeed still home. If the mobile node has just

returned from a foreign network, it must first de-register with its home agent by submitting a registration request message to its home agent.

If the mobile node determines that it is on a foreign network, it must obtain a care-of address on the foreign network. This care-of address can be either the address a foreign agent that has advertised itself on the foreign network, known as a *foreign agent care-of address*, or it can be a locally assigned address that has been set aside for use by visiting mobile nodes. This second type of care-of address is known as a *co-located care-of address*, because the address is actually assigned to the mobile node on the visited network. Co-located addresses are assigned through some external mechanism such as DHCP [Per96a].

When the mobile node is at a foreign network, it must register its new care-of address with its home agent. This is done by either sending a registration message directly to the home agent (in the case of a co-located care-of address) or by first sending the registration request to the foreign agent and then to the home agent (in the case of foreign agent care-of addresses).

When datagrams arrive on the home network for the mobile node, they are intercepted by the home agent and sent via an IP in IP tunnel to the visited network. The tunnel endpoint at the visited network will be either the mobile node itself (as with a co-located address) or the foreign agent. In either case, the tunnel endpoint must detunnel the datagrams and ensure they reach the mobile node. When the mobile node sends unicast packets while at the visited network, it sends packets to an available router on the

visiting network and packets are routed via conventional IP routing methods to their destination.

2.5 IP Mobility Support for Multicast

The above scenario describes how mobile IP operates for unicast datagrams. For multicast operation, RFC 2002 specifies two possible strategies.

To receive multicast datagram for a particular group, a mobile node must “join” the group using the IGMP messaging as described in Section 2.1.4. The mobile node has two options for joining a multicast group. First, if the visited network has an attached multicast router, the mobile node may join the multicast group locally. This is accomplished issuing an IGMP join request for the desired group with either its co-located care-of address or its home address as the source address for the IGMP join request. The mobile node may not use its foreign agent care-of address as the source of its IGMP messages [Per96a]. This method of joining is referred to as the remote subscription [HaW97] method.

The second method of joining a multicast group from a visited network is to create a bi-directional tunnel back to the home agent, which must also be a multicast router. The mobile node now tunnels its IGMP messages to the home agent and the home agent tunnels all datagrams for joined groups back to the mobile node [Per96a].

The home agent tunnels multicast datagrams to the mobile node according to the following rule. If the mobile node is using a co-located address, the home agent should tunnel the multicast datagrams directly to that address. Otherwise, the home agent must

first encapsulate multicast datagram inside a unicast datagram, and tunnel the unicast datagram to the mobile node. This adds a level of encapsulation allowing the foreign agent to know to which mobile node it must send the multicast datagram since the destination address of a multicast datagram is the group address. In either case, the mobile node must decapsulate the datagram it receives to extract the multicast datagram. The home agent can determine whether the mobile node is using a co-located or foreign host address by bits that are set in the mobile IP registration message that the mobile node sends to the home agent [Per96a].

Mobile nodes wishing to send datagrams to a multicast group also have two options for sending these datagrams. They may send multicast datagrams directly on the visited network or they can send them via a tunnel to the home agent. If sending directly on the visited network, the mobile node must use a co-located address as the source address for the multicast datagram since multicast routing protocols generally use the source address to make routing decisions. Similarly, if using a tunnel back to the home agent, the mobile agent must use its home address as the multicast source address.

2.6 The Mobile Multicast (MoM) Approach To Supporting Mobile IP Multicast

2.6.1 MoM Description

Harrison, et al [HaW97] have proposed a slightly different approach to handling mobile IP multicast. They introduce their own Mobile Multicast protocol for providing multicast support to mobile hosts. They claim that the current mobile IP model deals primarily with unicast routing issues, and that current multicast routing protocols

“implicitly assume static host when setting up multicast delivery trees” [HaW97]. Using either of the previously presented (RFC 2002) two options for mobile multicast would require tree reconstruction (costly) or intact use of trees with tunneling (leads to inefficiencies).

Remote subscription is well suited for hosts that spend a relatively long time at one particular visited network, since delivery trees would not have to be reconstructed very frequently. For mobile nodes can to use remote subscription, they must, either only receive multicast messages, or have a co-located address in order to be able to send datagrams to multicast groups [HaW97]. Also, the visited network must have a multicast router.

While bi-directional tunneling handles both recipient and source mobility, it may not provide the most efficient routing path. Consider a group consisting of two hosts from the same home network on both visiting the same distant foreign network. Bi-directional tunneling does not scale well either. This is because home agents must make as many copies of datagrams as they have away mobile nodes.

MoM was proposed to deal with the above-listed shortcomings of the IETF solution to mobile IP multicast. MoM put the burden of dealing with host mobility on mobility agents instead of on the multicast routers. Mobility agents expect their hosts to move while multicast routers generally do not. To handle node mobility, MoM utilizes foreign agents (FAs) in a slightly different manner than the IETF proposal for handling mobile IP multicast.

First, the MoM protocol does not allow remote subscription. Instead, all multicast traffic is routed to an FA and the FA uses local link layer multicasting to send multicast to mobile hosts on its subnet. This approach means that a multicast router on the visited network could, in fact, retransmit the multicast message that was transmitted by the foreign agent. This could in turn lead to unwanted multicast routing loops. To avoid this possibility, MoM specifies that all multicast messages delivered by an FA should have the TTL field set to one to assure no further retransmission beyond the local subnet.

The MoM approach avoids unnecessary duplication of messages by HAs, since an HA needs only to send one multicast message for all of its mobile agents on a given foreign subnet. One problem, however, that arises with handling mobile multicast in this manner is known as the “tunnel convergence” problem [ChW98]. Tunnel convergence occurs when many HAs have mobile hosts on a given foreign subnet subscribed to a common multicast group. In this case, the FA as well as the foreign network would be overwhelmed by duplicate multicast datagrams.

To eliminate this tunnel convergence, MoM introduces the idea of a Designated Multicast Service Provider (DMSP). A DMSP is simply the HA that a given FA chooses as the sole multicast provider for a particular group. Several DMSP selection policies (how the FA chooses the DMSP from available HAs) were studied in the development of MoM. Two of the best policies turned out to be the Oldest HA and Closest HA. The Oldest HA simply chooses the HA that has had mobile nodes on a given foreign network for the longest time, and the Closest HA policy simply chooses the HA that is topologically closest to the FA.

When a mobile host arrives at a foreign network, it registers with the FA and tell it to which multicast groups it is subscribed. The FA in turn registers with the mobile host's HA. This causes the HA to join the multicast group on behalf of the mobile host if not already joined. If the FA wishes to nominate the HA as a DMSP, it must also send a message stating so to the HA. Only upon receipt of the DMSP nomination message, will the HA forward multicast messages for the given group to the FA. When the last mobile host from a given HA leaves the foreign network, the HA should inform the FA of the fact that it will no longer be acting as a DMSP. The FA will then chose a new DMSP from among the remaining HAs servicing mobile hosts on the foreign network. The next step is to inform the new DMSP of its status. When the FA starts receiving datagrams from the newly selected DMSP, it releases the original DMSP from its DMSP responsibilities.

MoM requires the addition of several data structures to both HAs and FAs. HAs will have an away list to keep track of mobile hosts that are away plus other mobility binding information such as the FA address and the binding expiration time.

Both home HAs and FAs must keep track of group membership information on a per-group basis. HAs must track which mobile hosts are members of a particular group and at which FAs these groups reside. The HAs must also maintain a list of FAs for which they provide DMSP services. FAs also maintain a listing of all visiting hosts that are members of a particular group, and to which HAs these visiting hosts belong. FAs must maintain a list of which HA(s) are currently providing DMSP services for a particular

group. There may be more than one HA providing DMSP services if high reliability is desired.

2.7 MoM Simulation and Results

2.7.1.1 Simulation Model

MoM simulation model details are available in [WiH98]; important design aspects are presented here.

Williamson et al. simulated MoM operation using a discrete-event simulator. In the MoM simulation model, N uniformly distributed LANs were represented as points in a two-dimensional Cartesian coordinate system. Each LAN was home to H mobile hosts. There were G multicast groups with membership chosen at random from the available pool of mobile hosts. The number of hosts per group was set between 1 and 50 hosts per group. All mobile hosts started the simulation at their home network, and moved to randomly chosen destination nodes. Once a mobile host visited a foreign LAN, it returned home with probability of 0.5; otherwise it would chose another randomly selected foreign LAN to visit. Transit time, regardless of distance, was chosen to be constant. Each mobile host would spend an exponentially distributed amount of time at a foreign LAN before moving. The visit time value was set so that mobile host would spend 9.1% of time in transit, and 90.9% of their time connected to a LAN with 60.6% spent at foreign networks and 30.3% spent at home.

Exactly one stationary server (not located on any of the LANs) served each multicast group. The network topology interconnecting the LANs was not modeled, but the

Euclidean distance between sources and LANs was used to calculate routing efficiency. Routing distance was simply determined as the distance from the source to the DMSP LAN plus the distance from the DMSP LAN to the FA LAN. Routing efficiency was the ratio of this routing distance to the remote subscription distance. The remote subscription distance was taken to be the Euclidean distance directly from the source to the FA LAN.

2.7.1.2 Simulation Results

The creators of MoM set out to verify that their solution scaled well with respect to increasing group size, number of LANs and hosts, and number of multicast groups. They tested the handoff rates and routing efficiency and fairness of various DMSP selection policies. Finally, they looked at deliverability of multicast messages and overhead caused by MoM.

For determining scalability, the following statistics were collected: number of multicast group members per LAN, number of mobile hosts away (per HA), number of multicast group members away (per HA), number of foreign LANs currently visited (per HA), and number of current DMSP responsibilities (per HA). The key consideration for scalability was the load on the HAs. For MoM, the load on the HAs was directly related to the number of DMSP responsibilities. The MoM team conceded that, as far as multicast delivery is concerned, the IETF remote subscription method put the least load on the HA. This is because the HA does not need to concern itself with multicast traffic destined to any of its remotely subscribed mobile hosts. However, the load caused by the IETF bi-directional tunneling was directly related to the number of multicast group

members away per HA, and was shown to scale considerably worse than the MoM solution [WiH98].

The MoM team tested various DMSP selection algorithms including Newest MH, Random, Newest HA, Count Based, Closest to FA, Closest to Source, Oldest MH, and Oldest HA. Oldest HA yielded the least number of DMSP handoffs across all sizes of groups tested while Newest MH yielded the most number of DMSP handoffs. However, the Closest to FA method of DMSP selection method yielded the best routing efficiency. This approach requires, on average, a route of only 2 to 2.2 times that of remote subscription. The remaining policies average routes about 2.5 times that of remote subscription.

DMSP selection method fairness was determined by the average number of DMSP responsibilities associated with a given HA. A fair was considered as one that has all HAs sharing roughly equal numbers of DMSP responsibilities. It was determined that all selection algorithms produce fair results with the exception of location based algorithms. This was attributed to the fact that all FAs choose the same HA in the Closest to Source algorithm. A given FA will always choose one of its close neighbors in the Closest to FA algorithm.

The deliverability of multicast messages was shown to be adversely affected for multicast groups of size 5-20 because of incorrect DMSP state information. The incorrect DMSP state information resulted from a faulty DMSP handoff algorithm. The algorithm allowed for a DMSP to stop transmitting when its last mobile host left a LAN and registered at another LAN. The authors pointed out that this could have been alleviated if the FA was required to continue to serve as DMSP until handoff was

completed. Even so, the authors also demonstrated that algorithm performance could be improved by shortening the DMSP state timeout period.

The final evaluated aspect was the overhead that would be caused by DMSP selection messages. The team noted that a portion of DMSP selection messages could be piggybacked with HA registration messages. Messages can be piggybacked when a DMSP message is triggered by the arrival of an MH assigned to the newly selected DMSP. DMSP messages triggered by handoffs that occur asynchronously with the arrival of an MH from the new DMSP cannot be piggybacked. They showed that for small group sizes (< 5 hosts), most DMSP selection messages could be piggybacked, while for large group sizes (> 30 hosts), about one-tenth of the messages could be piggybacked. With large group sizes and efficient handoff algorithms, the effects of not being able to be piggyback are reduced because the number of DMSP messages sent gets proportionately much smaller than the number of MH registration messages.

2.8 Summary

This section reviewed the current state of IP multicast, including its management protocol (IGMP), routing protocols, and current implementations. Protocol dependent and protocol independent implementations of both shared tree and source-based algorithms were introduced. Next, IP in IP tunneling was discussed as a prelude to IP mobility support. IP Mobility support was introduced and the proposed IETF multicast support methodologies were presented. Finally, an alternative to the IETF proposals, MoM, was presented, along with results of performance analyses conducted by the MoM development team.

Chapter 3: Methodology

3.1 Introduction

This chapter presents the methodology used to develop and analyze the simulation model. Section 3.2 presents the problem, defines its scope, and justifies the choice of simulation for this research. Section 3.3 presents a new approach to mobile multicast transmission. Section 3.4 defines the operational assumptions made in creating the simulation model. Section 3.5 details simulation model design and operation. Section 3.6 discusses the different mobility support mechanisms to be compared, and Section 3.7 goes on to detail the experimental factors varied during the simulation runs. Section 3.8 defines the performance metrics used to compare the different mobility support mechanisms, and Section 3.9 details how simulation models were verified and validated. Finally, the chapter is summarized in section 3.10.

3.2 Problem Overview

As previously stated, little published research exists in the area of determining how currently proposed IP mobility support mechanisms affect the performance of IP multicasting. Further, the research that does exist in this area [WiH98], abstracts away the supporting network topology and routing mechanisms that would, in reality, support such a system. The research also overlooks key factors such as support for multiple multicast sources, and allowing the mobile hosts themselves to be sources.

3.2.1 Problem Definition

The literature review indicates that there is little research into the area of IP mobility support for IP multicasting. No information is available comparing the performance of currently defined mobility solutions for IP multicast under actual network conditions. The MoM solution presented in [HaW97, ChW98, and WiH98] is an alternative to the IETF proposals for allowing mobile host to join multicast groups. It does not, however offer any improvements on the mechanisms used to allow mobile hosts to act as multicast sources. To assess the performance of mobility support solutions for IP multicasting, both multicast join and transmission mechanisms must be analyzed.

3.2.2 Problem Statement

The focus of this research is to perform a comparative analysis of currently defined mobility solutions for IP multicasting. This research considers combinations of currently defined join and transmission mechanisms. It also introduces and compares a novel mobile multicast transmission mechanism.

3.2.3 Problem Scope

To solve this problem in a reasonable amount of time, it is necessary to limit the problem scope. Scope limitations applied to the aspects of network topology and dimensions, number of mobile nodes, number of multicast groups, number of multicast group members, and the number of multicast sources are detailed in the following subsections.

3.2.3.1 Network Topology and Dimensions

A representative military network topology is desirable for study. The Framework Nations Network (FNN) [Wen98] topology utilized by the Implementation Force (IFOR) in the Bosnian theater was chosen as the basis for the representative topology for this research. While the representative network topology does not precisely match the FNN topology, its physical scale provides a realistic backdrop for modeling a network infrastructure to support mobile nodes. Bosnian theater topographical features are not modeled. The rough topology is depicted in Figure 2.




Figure 2: Framework Nations Network Topology

3.2.3.2 : Number of Multicast Groups

Since multicast routers deal with multicast groups independently [WiH98], only one group will be analyzed. This may cause some lack of generality because interaction among groups cannot be observed, but the critical mechanism that support mobile multicast will still be observed.

3.2.3.3 Mobile Nodes

For this experiment the minimum number of mobile nodes is 40. This minimum was chosen because there are 40 fixed mobility agents, and with 40 randomly placed mobile nodes a dense multicast environment is assured. Scenarios with 80, and 120 mobile nodes are also presented to represent medium and high levels of loading.

3.2.4 Method of Evaluation

According to Jain, there are three possible methods of analyzing a system's performance: analytical modeling, simulation, and measurement [Jai91]. As previously mentioned, none the IP Mobility support mechanisms for multicast have been widely fielded, and consequently very few actual mobile IP networks exist that use IP mobility to support IP multicast. Furthermore, the FNN is unavailable for applying the technologies presented in this research. Even if this network were available, and the technology could be fielded within the FNN environment, the limited time allotted to this research is not sufficient to provide conclusive evidence that an improvement was the result of a parameter setting rather than a random change in the environment [Jai91]. For these reasons, direct measurement techniques were not chosen.

Analytical modeling techniques provide low accuracy because of low model fidelity due to the many simplifying assumptions that must be made to obtain results [Jai91]. Consequently, this method of evaluation was ruled out.

Simulation was chosen as the technique to evaluate the performance of four mobility support mechanisms for IP multicast. Simulation allows flexibility in the level of model detail. Simulation models can be validated to ensure that assumptions are reasonable, and, when correctly implemented, can produce behaviors and performance with a high level of fidelity to real world systems [Jai91].

3.2.5 Simulation Tool

The *Optimized Network Engineering Tools* (OPNET) *Modeler* was chosen as the modeling and simulation tool for this research. OPNET Modeler is a discrete-event simulator that allows hierarchical object-based modeling of networks and their component systems. OPNET has a wide variety of predefined network node and link models as well as an extensive set of tools for creating new node models or customizing existing models.

OPNET node models are comprised of processor modules that are defined using a combination of graphical finite-state machines and C or C++ code. The use of finite state machine models simplifies the definition of protocols and other processor interactions.

3.3 Minimal Multicast Encapsulation

The literature review presented two modes of mobile multicast transmission: direct transmission, and home tunneling. Direct transmission entails a mobile node transmitting multicast packets with its care-of address as the multicast source address.

Home tunneling entails encapsulation of multicast packets inside unicast packets in order to tunnel the multicast packet to the home network for decapsulation and distribution.

If direct transmission is used with co-located addresses, it provides correct multicast routing, and a unique but temporary return address in the source address. If this method is used with FA care-of addresses, no exact return address is available to recipients of multicast packets. With either choice of care-of address, recipients cannot discern the true multicast source identity from the source field. Since this research deals solely with FA care-of addresses, direct transmission of multicast packets on visited networks is not considered in this research. The home tunneling method described above is presented as the only alternative solution to the multicast transmission problem when using FA care-of addresses [Per96a, ChW98]. This method multicast of transmission appears to lead to possible routing inefficiencies and increased congestion at home agents that have many away mobile nodes acting as multicast sources. Therefore, to eliminate inefficient tunneling of multicast transmissions and allow use of FA care-of addresses while providing source identification, this research presents a novel mobile multicast transmission mechanism that will hereto forth be referred to as *minimal multicast encapsulation*.

Minimal multicast encapsulation is similar to unicast tunneling in that multicast IP datagrams are encapsulated in outer IP datagrams. They differ in the fact that, for minimal multicast encapsulation, the outer datagram is also a multicast datagram. The purpose encapsulation in minimal multicast encapsulation is to ensure correct multicast routing throughout the network while preserving source identity. To receive minimal multicast encapsulated packets, multicast hosts need to be able to recognize that the

minimal multicast encapsulated packets are indeed encapsulated and then un-encapsulate them. This can be accomplished in much the same manner as traditional minimal IP encapsulation [Per96c] with a few changes to account for the multicast destination vice a unicast destination.

As in minimal IP encapsulation, the minimal multicast encapsulation host inserts a minimal forwarding header between the original IP header and the data payload. The original header's total length field is incremented by 8 to account for the minimal forwarding header. The protocol field is set to 56 indicating minimal multicast encapsulation, and the original header's source address is changed to the mobile host's care-of address. All other original header fields remain the same except the header checksum, which is now recalculated to account for the change in field values.

The inserted minimal forwarding header protocol field is now set to the outer header's original protocol value, and its source address is set to the outer header's original source address. The minimal header checksum is now calculated over the minimal header fields. These headers are shown in Figures 3 and 4.

When the minimal multicast encapsulated packet is routed through the network, it will be correctly routed by source-based multicast routing protocols since the outer source is correct for the given network topology. When the minimal multicast encapsulated packet arrives at the destination, the receiving host recognizes the packet as being minimally multicast encapsulated by reading the protocol field. Upon receiving the minimal multicast encapsulated datagram, the receiver first verifies that all checksums are valid and then restores original header values from the values stored in the minimal header. Finally the receiver removes and discards the minimal header.

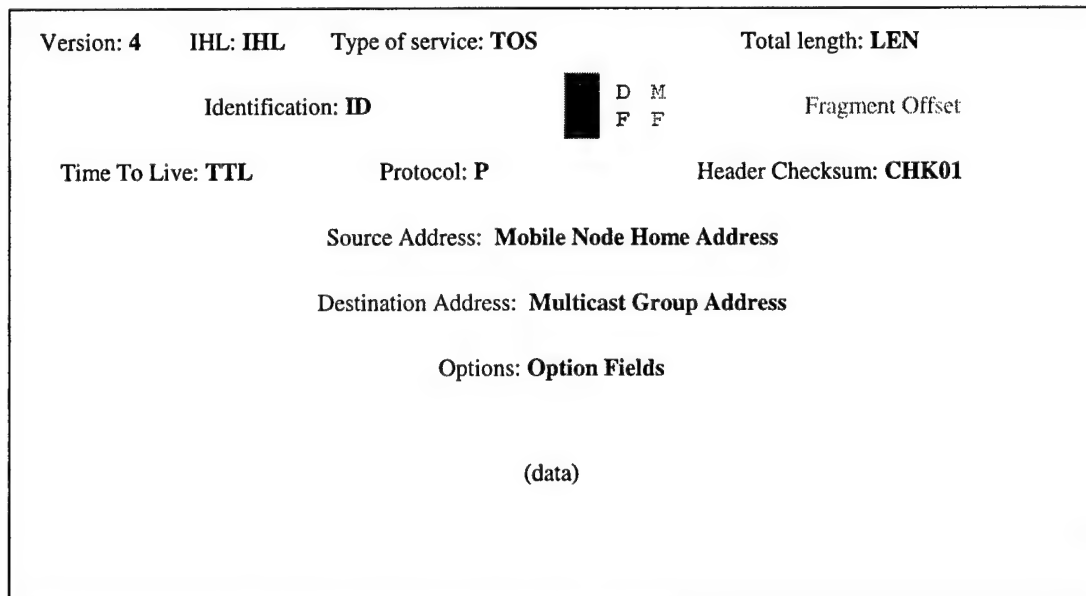
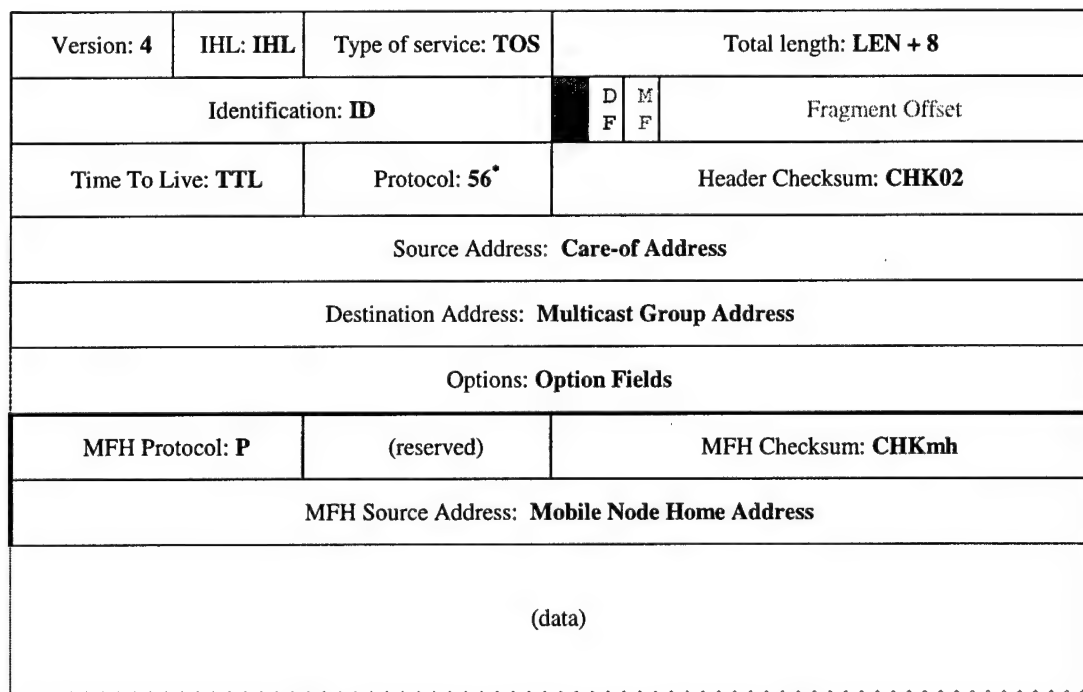


Figure 3: Original IP Header



*Protocol 56 is currently an unassigned protocol number

Figure 4: Minimal Multicast Encapsulated Header

3.4 Operational Assumptions

In the absence of a fielded mobile IP multicast system, several simplifying assumptions were made in the development of this model. As stated in the previous Scope subsection, some simplifying assumptions were also made to ensure timely completion of scenario runs. These assumptions are listed in the following subsections.

3.4.1 Mobile IP Configuration Assumptions

3.4.1.1 Care-of address Assignment

The IETF IP mobility support RFC [Per96a] specifies two methods for assigning care-of addresses to mobile IP nodes: co-located care-of addresses and foreign agent care-of addresses. Co-located care-of addresses require setting aside enough extra IP addresses to support some finite number of visiting hosts. This can lead to inefficiencies in utilization of a possibly limited resource, namely, the visited network's IP address space. Consequently, this research focuses solely on foreign agent care-of addresses.

3.4.1.2 Agent Advertisements

According to Perkins [Per96a], IP mobility agents can be configured to transmit advertisements either on a periodic basis or in response to solicitations. The later method of advertisement can be utilized if mobile hosts are guaranteed to transmit agent solicitations upon changing links [Per96a]. Since the link model used in this research is capable of detecting a link change, agents only advertise in response to solicitations.

3.4.2 Network Components

3.4.2.1 Routers

For this research, the interconnected nodes in the Framework Nations Network topology are considered to be high-speed routers (capable of routing 120,000 packets per second) with multicast routing capabilities. High-speed routers are used to remove possible network bottlenecks.

3.4.2.2 Mobility Agents

In addition to interconnecting links, each router is connected to a local IP mobility agent that utilizes a broadcast link to communicate with mobile hosts in its vicinity. Mobility agents serve as both home and foreign mobility agents for their areas of responsibility.

3.4.2.3 Mobile Nodes

For unicast IP transmissions, mobile nodes communicate solely with mobility agents via broadcast links. For multicast and broadcast IP transmissions, mobile nodes may be configured to communicate with other mobile nodes on the same link. This choice depends upon the version of IP mobility support for multicast is chosen for the current run.

3.4.3 Network Links

3.4.3.1 Fixed Link

Fixed links between routers are assumed to be duplex DS3 speed links. As with the routers, the DS3 links were chosen to remove network-induced bottlenecks.

3.4.3.2 Mobile Links

Since the purpose of this research is to evaluate IP-layer mobility support for multicast, the mobile link layer has been abstracted to an *ideal* wireless link. In this case, an *ideal* wireless link is one in which packets transmitted from a mobile link are only received by the nearest base station (mobility agent) and not received by more distant stations. Furthermore, all packets transmitted by a given base station are received only by those mobile nodes that are currently closer to the transmitting base station than to any other base station.

As mentioned above, this mobility link layer can be selectively configured to allow all mobile nodes that are currently in communication with a base station to receive packets sent by any other mobile node currently in communication with the base station.

3.4.4 Group Sources

Sources for the test multicast group are eight randomly selected mobile nodes transmitting at an application-layer rate of 8 kbps. Eight nodes were chosen to give a total application-layer throughput of 64 kbps (toll quality voice).

3.4.5 Group Membership

For a given simulation trial, the number of mobile group members equals the number of mobile nodes.

3.4.6 Multicast Routing Algorithm Selection

Given the representative network topology parameters and the fact that each node is expected to periodically have a member of each multicast group visiting, a dense mode multicasting algorithm is used. The multicast routing algorithm used is a variation on PIM-DM. As with PIM-DM, the algorithm employs prune and graft messages between routers to limit or expand the multicast distribution tree.

3.4.7 Application Arrival Rate and Size

The primary multicast data transmitted is multicast voice traffic. A representative coder-decoder (CODEC) for voice applications compresses each 20 ms of 8000 Hz, 16-bit sampled input speech into 266 bit packets [Mck99]. This sampling rate yields a mean arrival rate of 50 packets per second.

3.4.8 Background IP Traffic

To minimize simulation run times, background traffic is not included.

3.4.9 Mobile Node Movement

Mobile nodes choose random destinations from the set of mobility agent positions. It is assumed the network routers, and thus the mobility agents, are located in cities to where mobile nodes may wish to travel. A reason for this simplification is that mobility agent coordinates can be easily determined. Once a node chooses a destination, it travels

to there along a straight line. After the node reaches its destination, a new destination is chosen at random from all available mobility agents. Due to the mobile link nature, the node establishes links with the closest mobility agents along its path to the destination.

3.5 Model Design and Operation

3.5.1 OPNET IP Node Modules

OPNET Modeler provides predefined node models of various IP nodes including routers, workstation nodes, and server nodes. Variations on these nodes are provided that utilize point-to-point links, Ethernet links, and/or combinations of the two.

3.5.1.1 OPNET IP Process Modules

All OPNET IP nodes share a core set of processes to provide IP-based services and have one or more link transceivers. Depending upon the link to be modeled, nodes may have link-layer processes, such as address resolution or media access control, or they may have no additional link layer processes. The link-layer processes – or the transceiver processes in the case of point-to-point links – are then connected to the IP routing process.

The IP routing process determines necessary network-layer routing and fragmentation that must be performed for packets received from both upper and lower protocol layers. It then sends packets out on the appropriate interface according to the selected routing protocol.

The IP encapsulation process is above the IP routing process. This process encapsulates transport-layer packets in IP packets and passes them down to the routing

process. Conversely, it un-encapsulates packets received from the lower layers that the IP routing process has determined are destined for upper transport layers. The IP encapsulation process reads the protocol field to determine which transport-layer protocol process should receive the packets and then sends them to the appropriate transport protocol process.

The primary transport processes contained in all OPNET IP node models are the UDP process and the TCP process. These processes either encapsulate packets received from the upper-layer application processes in transport packets or un-encapsulate transport packets received from the IP encapsulation process.

OPNET provides a set of standard application-layer processes for workstation and server nodes, as well as, routing application processes for router nodes. The model designer can add applications that attach above either the TCP or UDP transport protocol processes. The abovementioned processes are contained in the OPNET-provided IP Router, Point-to-point Client, and Point-to-point Server node models depicted in Figures 5 and 6.

3.5.1.2 Changes to Existing OPNET IP Process Modules

The following subsections detail changes to the OPNET IP routing processes and sub-processes.




Figure 5: OPNET IP Router node model

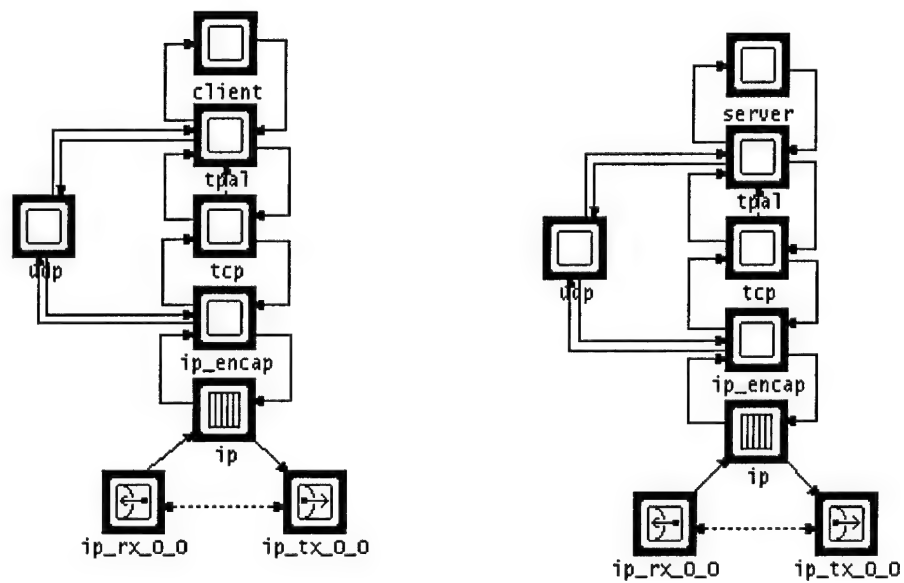


Figure 6: OPNET Point-to-point Client and Server node models

3.5.1.2.1 IGMP Processes

A newly updated (October 1996) OPNET IP process model provides three IGMP child processes for supporting both IP routers and hosts. First, the IGMP host process allows non-router hosts to join multicast groups by sending, receiving and processing the appropriate IGMP messages. Second, the IGMP Router Interface process serves as the interface between the IP Routing process and the router's multicast-enabled ports. One IGMP Router Interface instance process is created upon router initialization for each multicast enabled interface. Third, each IGMP Router Interface process creates one IGMP Router Group process instance for each subscribed group. The IGMP Router Group process maintains the state of a particular group on a given port. The IGMP Router Group process also communicates, as necessary, with the router's multicast routing process.

The updated OPNET IP process model does not yet support IP mobility; therefore, the IGMP host process was modified to allow for rejoining of a group upon link change as required for the remote subscription method of multicast reception. Additionally, the IGMP router processes were also modified. These modifications allow mobility agent routing processes to track away-host group memberships.

3.5.1.2.2 Multicast Routing Process

The updated OPNET IP Routing process model provides a PIM-SM multicast routing child process. OPNET does not provide a dense mode routing process. To provide dense-mode multicast routing, the PIM-SM process model was used as a framework for creating a PIM-DM-like process model. The basic OPNET PIM-SM data structures,

finite-state machine (FSM), message formats, and inter-process interfaces were retained. Additional data structures and FSM states were added, and operational code within the FSM states was extensively modified to provide dense mode and mobile operation.

3.5.2 New IP Node Models

Current OPNET nodes do not model IP mobility support in any form; neither mobile IP nodes nor mobility agent nodes currently exist. New IP mobility node models were created by adding new process modules to current OPNET IP node models. The remainder of this section provides a general description of how these new nodes were created.

3.5.2.1 Mobile Node Node Model

The mobile node model (Figure 7) for this research is based on OPNET's point-to-point workstation model. To provide IP mobility support to mobile nodes, the following process modules have been added to the mobile node model: mobile IP messaging, mobility support, and IP tunnel endpoint. To provide for minimal multicast encapsulation and decapsulation of multicast packets, a multi-tunnel encap-decap process was added. Mobile transceiver, link filter, and link router processes have been added to provide the *ideal* link functionality described earlier in this chapter. Finally a movement process was provided to give the node the random movement characteristics described previously. The following subsections detail the operation of these added modules.

3.5.2.2 IP Mobility Support Modules

The three IP mobility support process modules work in conjunction to provide IP mobility support described in [Per96a] and in Section 3.3.

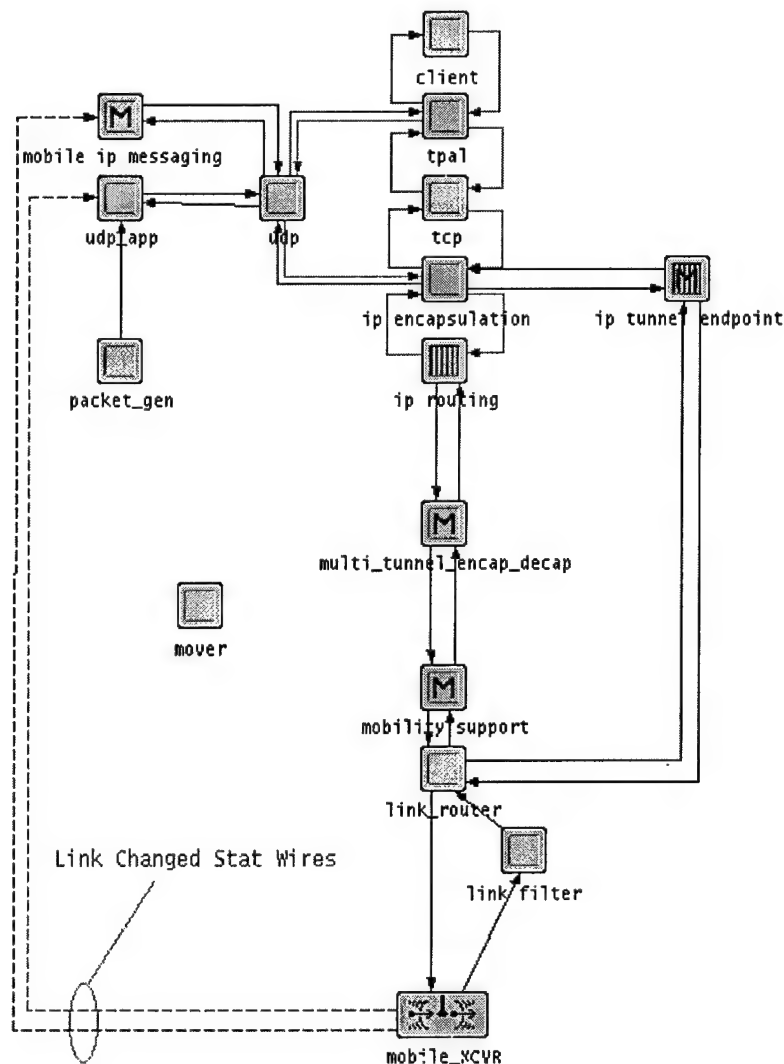


Figure 7: Mobile Node Node Model

3.5.2.2.1 Mobile IP Messaging Process Module

The mobile IP messaging model provides a UDP-based messaging facility for the mobile node to send registration requests and receive registration replies. As described in [Per96a], registration requests and replies are sent to UDP port 434.

3.5.2.2.2 Mobility Support Process

This process is located between the IP and link layer processes and is responsible for forwarding packets to the link router process module. Along with the packet, this process also sends link layer routing requests via an Interface Control Information (ICI) data structure. This ICI tells the link router whether to tunnel the IP, to send it out directly, or to send it directly and request tunneling on the link's distant end. The ICI also has fields that can request that the distant end receiver rebroadcast the link layer packet so that other mobile nodes on the link can receive the sent packet. This process determines ICI field settings by accessing a combination of information in the Mobile IP Messaging Process's mobility support tables and its current IP mobility support settings.

3.5.2.2.3 IP Tunnel Endpoint Process

The IP tunnel endpoint process provides IP encapsulation of packets received from the link router by passing them directly to the IP encapsulation process along with an ICI indicating the IP tunnel endpoint addresses. As stated above, this information is determined by examining the mobility support information maintained by the mobile IP messaging process.

3.5.2.2.4 Multi-Tunnel Encapsulation-Decapsulation Process

The multi-tunnel encap-decap process provides minimal multicast encapsulation and decapsulation of multicast packets sent by mobile nodes using the minimal multicast encapsulation transmission mechanism as described in Section 3.3. A separate process was created to handle minimal multicast encapsulation, as the mechanism utilizes minimal encapsulation and does not require re-routing through the IP encapsulation process.

3.5.2.3 Mobile Node Link Processes

To provide the *ideal* wireless link functionality for the mobile node, two process modules were created: the link router and the mobile transceiver.

3.5.2.3.1 Link Router

The link router provides link layer encapsulation of upper layer packets to be sent to the transceiver and decapsulation of link layer packets to be sent to the upper layer. The link layer packet provides the ability to choose a link destination at the far end for transmitted packets and informs the link router of the near-end destination of received packets. The destination field, shown in Figure 8, can be set to select one of two upper layer destinations: the IP layer or the IP tunnel endpoint.

A special flag is also present in the link layer packet to indicate the sending mobile node is requesting link-layer rebroadcast. If this flag is set, the receiving base station link router not only sends the inner IP packet to the appropriate upper layer process (the IP processes or IP tunnel endpoint), but it also sends a copy of the entire link layer packet

back to the transceiver for broadcast to all mobile receivers currently communicating with the base station.

Far End Link Destination 1 = IP Processes 2 = IP Tunnel Endpoint
Link Broadcast Requested 0 = no 1 = yes
Data (upper layer packet)

Figure 8: Link Packet Fields

The link router also provides routing of upper layer packets based upon the link layer ICI depicted in Figure 9. The link layer ICI has a near end destination field that indicates how packets from the upper layer are to be routed. This allows both the IP processes and the Tunnel Endpoint to send un-encapsulated IP packets to each other or link encapsulated IP packets to the transceiver. Upper layer processes also use the link layer ICI to indicate how to set the fields in the link layer packet header for packets that are routed to the transceiver.

Near End Link Destination 0 = Link Transceiver 1 = IP Processes 2 = IP Tunnel Endpoint
Far End Link Destination (set only if Near End Dest =0) 1 = IP Processes 2 = IP Tunnel Endpoint
Link Broadcast Requested (set only if Near End Dest =0) 0 = no 1 = yes

Figure 9: Link ICI Fields

3.5.2.3.2 Link Filter

The Link Filter process emulates link-layer addressing protocols by accepting only unicast packets whose IP address matches the mobile node's IP interface address. The link filter process also emulates link-layer multicast reception by only accepting multicast packets addressed to groups to which a mobile node is subscribed.

3.5.2.3.3 Mobile Transceiver

The Mobile Transceiver process functions to transmit and receive link packets to and from the base station Transceiver process in the nearest mobility agent. All mobile transceiver processes utilize a global list of base station positions to determine which base station is closest. When link layer packets are received from the link router, the mobile transceiver calculates the distance to the base station and delivers the packet to the nearest base station. Upon determining the nearest base station transceiver, the mobile Transceiver adds its process identifier and coordinates the base station transceiver's list of mobile nodes that are in communication with it. In this way, the base stations know to which mobile nodes they must deliver packets. Additionally, upon change of nearest base station, the mobile transceiver sets the "link changed" local statistic to one. In this way, any upper-layer module that is connected to the mobile transceiver via a statistic wire will be notified of link changes.

The mobile transceiver process determines the nearest base station every ten seconds of simulation time. An overlap factor of 100 meters is included to prevent unnecessary switching back and forth between base stations when a mobile node is traveling close to the line that is equidistant to two base stations.

To send a packet, a mobile transceiver calculates the distance to the nearest base station and the associated propagation delay. It then sends the packet to the station by using the OPNET kernel procedure `op_pk_deliver_delayed`. In this way, the packet delivery is delayed by the propagation delay. When packets are received from a base station, the Mobile Transceiver simply forwards the packet to the link router.

3.5.2.4 IP Mobility Agent Node Model

The mobility agent node model shown in Figure 10 has IP and Transport layer processes identical to the mobile mode node model. The mobile agent's IP routing process gateway function is enabled so that the agent may act as a router. Additionally, application layer processes are replaced with upper layer processes required to implement routing protocols. In addition to its *ideal* wireless link, the mobility agent also has a point-point link for connection to the fixed network. While many mobility support and link layer processes in the mobility agent node model share names and basic functionality with the mobility support and link layer processes in the mobile node node model, there are some differences in exact functionality and structure. These differences are detailed in the following subsections.

3.5.2.5 IP Mobility Support Modules

As in the mobile node, the mobility agent's three IP mobility support process modules work in conjunction to provide IP mobility support described in [Per96a].

3.5.2.5.1 Mobile IP Messaging Process Module

This process provides the messaging service required for mobility agents to receive registration requests and either generate registration replies while acting as home agents or forward registration requests while acting as foreign agents. As implemented in the mobility nodes, these messaging services utilize UDP port 434. Additionally, this process maintains the home and visiting agent tables.

3.5.2.5.2 Mobility Support Process

This process has the same basic functionality as the mobile node mobility support processes except that its link layer routing decisions are based upon the mobile agent responsibilities for the currently selected IP mobility multicast support mechanisms.

3.5.2.5.3 IP Tunnel Endpoint Process

The IP Tunnel Endpoint Process is identical to the mobile node's IP tunnel endpoint process.

3.5.2.6 Link Processes

The link processes provide the base station side of the *ideal* wireless link.

3.5.2.6.1 Link Router

The link router is identical to the mobile node's link router.

3.5.2.6.2 Base Station Transceiver

The base station transceiver's primary function is to send and receive packets to and from mobile transceivers. As stated earlier, mobile transceivers only communicate with the current closest base station. To keep track of which mobile nodes consider the base

station their current closest node, each base station maintains a list of mobile transceiver processes that consider it to be the closest base station. As previously stated, these lists are actually populated and updated by the mobile transceivers themselves.

When a packet is received from the higher layer, the base station simply sends a copy to each mobile transceiver currently in its list. To do this, the base station calculates the distance to each mobile transceiver in the list and sends the copy of the packet delayed by the propagation delay.




Figure 10: IP Mobility Agent Node Model

3.5.2.6.3 Transmit and Receive Queues

To set the maximum data rate for the *ideal* wireless link, transmit and receive queues are used between the link router and the base station receiver. Since only one base station transceiver can send a packet to a given mobile transceiver, the transmit queue's processing rate sets the data rate for data transmitted by a base station transceiver. A transmit queue at in the mobile transceiver does not provide the same limit for traffic from a mobile transceiver to the base station transceiver because several mobile transceivers can send packets to a given base station at one time. A receive queue is therefore utilized to provide the maximum data rate limit for traffic coming from the mobile transceiver.

3.6 IP Mobility Multicast Support Mechanisms Tested

For this research, experiments are conducted on each of four possible combinations of multicast join and transmission mechanisms detailed in Table 6.

3.6.1 Multicast Join Mechanisms

This research examines the performance of the two IETF multicast join mechanisms discussed in the literature review: remote subscription and bi-directional tunneling.

3.6.2 Multicast Transmission

The two schemes considered are the home tunneling scheme recommended by the IETF and MoM in [Per96a] and [ChW98] respectively, and the novel minimal multicast encapsulation transmission mechanism introduced in Section 3.3.

3.6.3 Mobile Join and Transmission Combinations

This research compares each of four possible combinations of join and transmission mechanisms individually: i.e. all mobile nodes in a given scenario implement the same combination for the scenario duration.

Table 6: Mobile Join and Transmission Combinations

Join Mechanism	Transmission Mechanism
Remote Subscription	Home Tunneling
Remote Subscription	Minimal Multicast Encapsulation
Bi-directional Tunneling	Home Tunneling
Bi-directional Tunneling	Minimal Multicast Encapsulation

3.7 Experimental Factors

To test and compare the scalability, routing efficiency, and packet loss rate, of each transmit-receive mechanism combination, the number of mobile nodes and thus multicast mobile group size are varied.

3.7.1 Mobile Group Size

Section 3.4.5 states that the number of mobile multicast group members is equal to the number of mobile nodes, and section 3.2.2 states that the number of mobile nodes varies from 40 to 120 in increments of 40. Therefore, the group size varies accordingly.

3.8 Performance Metrics

To compare the four combinations of join and transport mechanisms, the following performance metrics are collected.

3.8.1 Average Relative Path Length

The relative path length is the ratio of actual path traveled by a multicast packet from source to destination divided by the shortest path length from source to destination. Path lengths are measured in hops as well as in physical distance. The average relative path length is simply the arithmetic average of relative path lengths for all multicast packets.

3.8.2 Number of Lost Packets

Several circumstances can lead to packet loss. Of primary interest are packets lost due to receiver link changes. Metrics for packet loss are collected and reported in terms of total number of packets lost, the number of packets lost due to link changes, and the ratios of number of receiver link changes that experience loss (degraded link changes) divided by the total number of receiver link changes.

3.9 Model Verification and Validation

The simulation model used in this research is verified and validated to ensure sufficient fidelity with proposed systems for providing mobility support for IP multicast. Verification ensures that the simulation executes correctly in accordance with model design and assumptions. Validation attempts to ensure that the assumptions made in designing the simulation model will, if properly implemented, produce results close to those observed in real systems [Jai91].

3.9.1 Model Verification

Verification entails ensuring that simulation code executes as designed. Several techniques were used to verify correct simulation model operation. First, a modular design

approach was taken. OPNET's subnet-node-module-process hierarchy simplified this approach. At the process level, the OPNET finite-state machine model aided in ensuring proper protocol behavior. Second, OPNET's built-in error detection was used to detect several instances of incorrectly designed state transitions. OPNET animations were used extensively to verify correct operation of protocols. The OPNET debugger was extensively used in combination with OPNET diagnostic blocks to discover and correct logical errors.

3.9.2 Model Validation

Law and Kenton suggest that a good technique to help ensure a valid model is to "collect high quality information about and data on the System" [LaK00]. Sources for gathering this information include consultation with subject matter experts (SMEs), observation of an existing system, review of existing theory, review of results from similar studies, and the experience and modeler intuition [LaK00]. The thesis advisor and thesis committee members, considered SMEs, were consulted regarding various design assumptions. Of the technologies studied in this research, stationary IP multicast is the only one that is currently commercially fielded. There is at least one academic research effort that has fielded a limited system to support IP mobility for unicast transmission [FID99], but there are no currently fielded systems that provide mobility support for IP multicast. As such, direct observation was not possible. The specifications for providing mobility support to IP multicast exist as proposed standards presented in the form of IETF RFCs and, in some cases, more informally as Internet Drafts. As such, models used in this research were validated against specifications put

forth in the abovementioned IETF documents, and available research, both considered existing theory.

Results from the model are also compared with theoretical results obtained from previous research efforts [WiH98]. The results of [WiH98] indicated that MoM had a path efficiency of between 2.2 and 2.8 while the results of this research indicate average path efficiencies of between 2.0 and 3.0.

Model validation was also aided by verification efforts in two instances. In the first case, verification tools helped to locate an invalid design assumption prompting a review of all available literature to correct the assumption. The design flaw was caused by incorrect assumptions made about the operation PIM-DM multicast routing. An Internet Draft that discussed PIM-DM [Dee99] made no mention of sending prune messages out on outgoing point-to-point interfaces that received incoming multicast packets. The document simply stated that these incorrectly received packets would be discarded and not forwarded. Upon reviewing OPNET generated animations used for verification purposes, it became clear that if these interfaces were not pruned, multicast packets would continue to be forwarded where they were not needed. This was because without sending these prune messages, only leaf nodes in the multicast delivery tree would be correctly pruned. A thorough literature review indicated that previous, expired, Internet Drafts detailed this aspect of pruning, but the current Internet Draft did not.

In a second case, verification tests led to the discovery of a design flaw in the OPNET provided IGMP Router Group model. During verifications testing of mobility support mechanisms for IP multicast, it was discovered that the OPNET IGMP Router Group process model could not handle multiple successive leave requests. Designers had made

the assumption that only that last host to have transmitted a group membership report would transmit a leave request. The IGMP RFC [Dee89] states that this is the preferred behavior, but it also states that hosts may transmit leave request out of turn if they lack the resources to know which host transmitted the last join request. When various hosts sent multiple leave requests during testing, the IGMP Router Group process crashed because it had entered into a state that had no valid transition to deal with a second leave request. This deficiency was easily corrected and proper operation verified.

3.10 Summary

In this chapter the problem was reviewed as well as, scoping issues, method of evaluation, and choice of simulation tool. The new mobile multicast transmission methodology, minimal multicast encapsulation, was introduced. Next operational assumptions were discussed. Model design and operation were then described to include changes made to existing OPNET models as well as new models that were created. The choices of mobility support mechanism to be tested were then describe. Next the experimental factors, and performance metrics to be collected were discussed. Finally, model validation and verification techniques were described.

Chapter 4: Results and Analysis

4.1 Introduction

This chapter provides analysis of simulation results. Section 4.2 discusses the statistical accuracy of presented results. Section 4.3 discusses the configuration and input factor settings for each scenario. Section 4.4 discusses the data collection methodologies for each performance metric. Section 4.5 presents an analysis and comparison of each performance metric for each combination of input and output support mechanisms described in chapter 3. Section 4.6 concludes with a brief summary of results.

4.2 Statistical Accuracy

Four combinations of mobile multicast transmission and reception mechanisms are presented and compared in this research. Each combination is tested at three distinct load levels. For this research, the number of mobile receivers determines the load level. All scenarios are independently replicated using four different random seeds; the random seed determines the starting locations of mobile nodes, and the paths that they will take. For a given seed and number of mobile nodes, the starting positions and paths taken are the same, regardless of the transmit or receive mechanism chosen. The Poisson transmission distribution, while determined by the seed, is not necessarily guaranteed to be the same from one combination of mechanisms to another, even if the seed and number of mobile nodes are identical. The choice of four different random seeds guarantees that performance metrics are not casually affected by changes in initial node positions, movement paths, or the Poisson traffic distribution.

Mean results from similar simulations using n different seeds are averaged to produce an overall mean of means, $\bar{X}(n)$. The variance of the means, $S^2(n)$, is calculated and a $100(1 - \alpha)$ percent confidence interval is calculated for each metric using the student's t-distribution. This confidence interval is given by Equation 1 [LaK00].

$$\bar{X}(n) \mp t_{n-1, 1-\alpha/2} \sqrt{\frac{S^2(n)}{n}} \quad (1)$$

For this research a 90 percent confidence interval was chosen, and $n = 4$ seed values were used

4.3 Simulation Scenarios

For research simulation scenarios, the factors varied included the transmit mechanism, the receive mechanism, and the number of mobile nodes. Simulation parameters included the number of mobile transmitters, the number of fixed transmitters, the data transmission rate, and the data transmission traffic distribution.

4.3.1 Simulation Execution Length

As previously stated in section 4.2, multiple independent repetitions are performed to establish confidence intervals on the means of measured performance metrics. To guarantee the results of multiple independent repetitions are not affected by start up conditions --as all nodes begin at their home network-- a warm up period is chosen in which nodes simply travel on their predetermined random path but do not transmit or receive multicast packets. At the end of the warm up period, multicast transmitters begin

transmitting to the test group. A warm up period of 10 hours was chosen to allow sufficiently random dispersion of mobile nodes throughout the network.

Once the warm up period was chosen, the transmission duration had to be chosen as well. Because of great variability of path efficiency, the degraded link changes to total link ratio was chosen as the performance metric to set the length of simulation scenario runs. As described in Section 3.8.2, the degraded link changes to total link changes ratio is the ratio of the number of receiver link changes that suffer packet loss due to link changes divided by the total number of link changes. The total number of link changes is equal to the number of time each receiver changes a link multiplied by the total number of transmitters.

As the above metric is a proportion, and no prior estimate is available, the total number, n , of samples (link changes) required to estimate the population proportion, p , is given by Equation 2 [MiA95].

$$n \doteq \frac{z_{\alpha/2}^2}{4d^2} \quad (2)$$

In Equation 2, d , is the maximum difference between the actual population proportion, p , and the sample proportion \hat{p} [MiA95]. Accordingly, to obtain a 90 percent confidence interval for p within one percent 6806 link changes are required. Trial runs determined that for scenarios with 40 mobile nodes 900 seconds of transmission time would be sufficient to allow for the required number of link changes. For the four seeds with 40 nodes, the number of link changes varied from 7104 to 7808. The number of link changes increases proportionally with the number of mobile nodes. Therefore, the transmission durations required to achieve an equivalent level of accuracy

for scenarios with 80 and 120 nodes were 450 and 300 seconds respectively. Trial runs verified that link changes varied from 7104 to 7808 for 80 nodes and from 7560 to 7808 for 120 nodes.

4.4 Data Collection Methodologies

This section details the methodologies and mechanisms employed to collect data for each performance metric.

4.4.1 Path Efficiency



Path efficiency is calculated in terms of hops and in terms of actual distance traveled by a packet. Application packets contain two zero-length fields that are used to store the number of hops taken and the distance traveled. Two additional zero-length fields contain the source node object identifier, and the source node's current mobility agent object identifier. As a packet passes from node to node the hop field is incremented by one and the distance field is incremented by the distance from the last node. When the application packet arrives at its final destination, the source node identifier is read along with source node's mobility agent identifier. These two numbers form an index into a two-dimensional array of records. Each record contains a sum of hop ratios field, a sum of distance ratios field, and number of ratios field. The shortest path (in both hops and distance) is determined and the path efficiency ratios for hops and distance are calculated. The path ratio is the length actual path taken divided by the shortest path length. Once the path ratios are calculated for a given packet, they are added to the indexed record's sum of ratios field, and the total ratios field is incremented. Then, when a receiver node changes links, the average path efficiency for each source-agent pair is calculated by

dividing the sums of ratios field by the total number of ratios. These two averages are then written to the hop ratio and path ratio global statistics, for each source-agent pair. After the statistics are written, the values in each record are reset to zero so that data collection can begin for the new link. The path ratio collection methodology is depicted in Figure 11.

By taking the average of path length ratios, this collection methodology results in the normalization of path lengths. Since this research is concerned with the routing efficiency of the support mechanism on any given combination of source-destination pairs, it takes the average path length for each combination observed. In this way if a mechanism produced an actual path of length 3 with an optimal path of 1 and an actual path of length 50 with an optimal path of 10, the average path ratio would be 4. This ratio simply tells the efficiency per link change, but does not give any information about what the actual overall incurred distance cost.

4.4.2 Packet Loss


Transmitters produce sequentially numbered packets starting with packet zero. Receivers store the identifier value of the last packet received from each transmitter. A receiver counts packets as lost, if it receives a packet with a packet number greater than one plus the last packet number received from a given source. A receiver considers all packets between the last packet received and the new packet to be lost. Packets received that have numbers less than the last packet received are not considered lost nor do they replace the last packet received.



Time = t1: mobile node 88 receives packet 1 transmitted by mobile node 41 while linked to fixed node 1

Optimal distance = $72+30+17+61 = 180\text{km}$
 Actual distance = $72+27+34+53+61 = 247\text{km}$
 Distance ratio = $247/180 = 1.37$
 Optimal hops = 4
 Actual hops = 5
 Hop ratio = $5/4 = 1.25$


Record index: 41,1
 Sum dist ratio: $0+1.37 = 1.37$
 Sum hop ratio: $0+1.25 = 1.25$
 Number of ratios: $0+1 = 1$



Time = t2: mobile node 88 receives packet 2 transmitted by mobile node 41 while linked to fixed node 1

Optimal distance = $65+30+17+39 = 151\text{km}$
 Actual distance = $65+27+34+53+39 = 218\text{km}$
 Distance ratio = $218/151 = 1.44$
 Optimal hops = 4
 Actual hops = 5
 Hop ratio = $5/4 = 1.25$


Record index: 41,1
 Sum dist ratio: $1.37+1.44 = 2.81$
 Sum hop ratio: $1.25+1.25 = 2.5$
 Number of ratios: $1+1 = 2$



Time = t3: mobile node 88 receives packet 3 transmitted by mobile node 41 while linked to fixed node 2

Optimal distance = $58+34+53+32 = 177\text{km}$
 Actual distance = $65+27+34+53+39 = 177\text{km}$
 Distance ratio = $177/177 = 1.0$
 Optimal hops = 4
 Actual hops = 4
 Hop ratio = $4/4 = 1.0$

Record index: 41,1
 Record index: 41,2
 Sum dist ratio: $0+1 = 1$
 Sum hop ratio: $0+1 = 1$
 Number of ratios: $0+1 = 1$



Time = t5:
 At this time no further packets have been received by mobile node 88 since time = t4, but mobile node 88 has changed links from node 5 to node 4.
 Since a link change has occurred, mobile node 88 writes the following statistics:

Record index 41,1 distance ratio = $2.81 / 2 = 1.405$ written to global distance ratio stat.
 Record index 41,1 hop ratio = $2.5 / 2 = 1.25$ written to global hop ratio stat.

Record index 41,2 distance ratio = $1.0 / 1 = 1.0$ written to global distance ratio stat.
 Record index 41,2 hop ratio = $1.0 / 1 = 1.0$ written to global hop ratio stat.

Figure 11: Path ratio collection methodology

Packet loss statistics are collected during at two distinct times: immediately after a receiver has changed links and between link changes. Of primary interest are the packets lost immediately after a receiver link change, because these losses are most likely related to latencies inherent in the given mobile multicast receive (join) mechanism. The number of link changes that suffer from packet loss are tallied and presented as a ratio of total number of degraded link changes to total number of link changes as described in section 4.3.1.

4.4.3 Required Throughput

Required throughput is a measure of the load on a given Base Station transmitter. This is equivalent to the throughput that a mobility agent, and the associated communications link, must be able to support to provide multicast mobility support to mobile nodes. To save output file space, the throughput statistic is written as the average number of bits per second that a transmitter transmits during a ten second interval. A transmitter simply counts the bits that it transmits during a given interval divides the total value by ten and writes the statistic.

4.5 Analysis and Comparison of Performance Metrics

This section presents, analyzes and compares results obtained for each combination of mobile multicast transmit and receive mechanisms presented in section 3.6.3.

4.5.1 Analysis of Path Ratios

Path ratios are reported in terms of hop ratios and distance ratios. The following subsections present the path ratio results for each combination of send and receive mechanisms.

4.5.1.1 Hop Ratio

Tables 7 through 10 present the hop ratio results for the four tested combinations of send and receive mechanisms. The combination that utilizes tunneling for both transmitting and receiving suffers from the largest average hop ratio (Table 7). The combination that utilizes no tunneling at all (Table 10), as expected, enjoys a hop ratio of 1. The two methodologies that utilize tunneling in only one direction have path ratios in between the best and worst cases.

Table 7: Hop ratio results for Bi-directional tunneling with Home Tunneling

Number of Mobile Nodes	90 % Confidence Interval Lower Bound	Mean	90 % Confidence Interval Upper Bound
40	2.864	3.091	3.317
80	2.968	3.056	3.143
120	2.931	3.203	3.474

Also of interest are the maximum observed hop ratios from the observed scenarios. These results are presented in Table 11. These maximums serve to illustrate the fact that while the two methods that tunnel only in one direction, may have average hop lengths between the best and worst case, they can still occasionally suffer from very high maximum path lengths as compared to the mechanism that does not utilize unicast tunneling.

Table 8: Hop ratio results for Bi-directional Tunneling with Minimal Multicast Encapsulation

Number of Mobile Nodes	90 % Confidence Interval Lower Bound	Mean	90 % Confidence Interval Upper Bound
40	2.062	2.156	2.249
80	2.040	2.112	2.154
120	2.062	2.128	2.194

Table 9: Hop ratio results for Remote Subscription with Home Tunneling

Number of Mobile Nodes	90 % Confidence Interval Lower Bound	Mean	90 % Confidence Interval Upper Bound
40	1.999	2.179	2.359
80	2.093	2.154	2.216
120	2.030	2.289	2.548

Table 10: Hop ratio results for Remote Subscription with Minimal Multicast Encapsulation

Number of Mobile Nodes	90 % Confidence Interval Lower Bound	Mean	90 % Confidence Interval Upper Bound
40	1.0	1.0	1.0
80	1.0	1.0	1.0
120	1.0	1.0	1.0

Table 11: Maximum Observed Hop Ratios

Mechanism Combination	Maximum Observed Hop Ratio
Bi-directional Tunnel with Home Tunnel	15
Bi-directional Tunnel with Minimal Multicast Encapsulation	11
Remote Subscription with Home Tunnel	11
Remote Subscription with Minimal Multicast Encapsulation	1

4.5.1.2 Distance Ratio

Tables 12 through 15 present the distance ratio results for the four tested combinations of send and receive mechanisms, and table 16 presents the maximums. As is evident from the table the distance ratios have a greater variability than the hop ratios. This is attributed to the fact that the metric can take on an infinite number of values while the hop ratio metric is taken from a finite number of possible combinations. The same basic observations made for hop ratios can be made for distance ratios. Of note is the fact that the means and the maximums are even greater than those for hop ratios. While the distance ratio for the non-tunneling methodology is expected to be one, it is indeed slightly greater. Upon further analysis of simulation runs, this was attributed to the fact that the unicast routing algorithm used, RIP, used the hop distance as its metric and not actual distance or delay. This means that the shortest path chosen by the routing algorithm may not indeed be the shortest physical path.

Table 12: Distance Ratio Results for Bi-directional tunneling with Home Tunneling

Number of Mobile Nodes	90 % Confidence Interval Lower Bound	Mean	90 % Confidence Interval Upper Bound
40	4.548	5.031	5.515
80	4.264	4.335	4.407
120	4.233	4.866	5.498

Table 13: Distance Ratio Results for Bi-directional Tunneling with Minimal Multicast Encapsulation

Number of Mobile Nodes	90 % Confidence Interval Lower Bound	Mean	90 % Confidence Interval Upper Bound
40	3.077	3.320	3.563
80	2.818	2.896	2.973
120	2.866	3.015	3.163

Table 14: Distance Ratio Results for Remote Subscription with Home Tunneling

Number of Mobile Nodes	90 % Confidence Interval Lower Bound	Mean	90 % Confidence Interval Upper Bound
40	3.190	3.548	3.906
80	2.865	3.008	3.152
120	2.870	3.435	4.010

Table 15: Distance Ratio Results for Remote Subscription with Minimal Multicast Encapsulation

Number of Mobile Nodes	90 % Confidence Interval Lower Bound	Mean	90 % Confidence Interval Upper Bound
40	1.038	1.050	1.062
80	1.037	1.043	1.047
120	1.017	1.036	1.055

Table 16: Maximum Observed Distance Ratios

Mechanism Combination	Maximum Observed Hop Ratio
Bi-directional Tunnel with Home Tunnel	178.32
Bi-directional Tunnel with Minimal Multicast Encapsulation	178.31
Remote Subscription with Home Tunnel	179.07
Remote Subscription with Minimal Multicast Encapsulation	1.55

4.5.2 Analysis of Packet Loss

Two statistics are of primary interest with regard to packet loss: the degraded link ratio and the link change packet loss ratio. As previously stated the degraded link ratio is the ratio of the number of links changes that suffer from packet loss divided by the total number of link changes. The link change packet loss ratio is the total number of packets lost due to link change to the total possible number of packet that could be received. The following subsections present packet loss results for each combination of transmit and receive mechanisms.

4.5.2.1 Degraded Link Change Proportion

Tables 17 through 20 present the 90 percent confidence intervals for mean degraded link change proportion at each load level tested. It appears that the single factor that affects this metric the most is the receive mechanism. The mean values for degraded link ratios are an order of magnitude greater for the two combinations that rely on bi-directional tunneling than for those that use remote subscription. This attributed to the fact the mobile receiver must rely on the home agent to tunnel multicast packets, and the home agent must wait until it receives a move notification from the mobile agent to begin tunneling to the new location. There is inherently a greater delay in this process than in the local join process used in remote subscription. This is because, upon link change, a node using bi-directional tunneling must first send a solicitation to the new mobility agent. It must then wait for an advertisement from the new mobility agent so that it can determine the new care-of address. Only after an advertisement is received can the mobile node send a registration message to its home agent. On the other hand, a node

that is utilizing remote subscription is free to transmit an IGMP join request on the new network as soon as it detects a link change. Another possible additional difference in delay may also be caused by the fact that, in a dense multicast environment, the nearest join point on the multicast tree is likely to be closer to the mobile node than its home agent.

Table 17: Degraded Link Change proportion Results for Bi-directional tunneling with Home Tunneling

Number of Mobile Nodes	90 % Confidence Interval Lower Bound	Mean	90 % Confidence Interval Upper Bound
40	0.0129	0.0149	0.0168
80	0.0120	0.0149	0.0178
120	0.0131	0.0144	0.0157

Table 18: Degraded Link Change proportion Results for Bi-directional Tunneling with Minimal Multicast Encapsulation

Number of Mobile Nodes	90 % Confidence Interval Lower Bound	Mean	90 % Confidence Interval Upper Bound
40	0.0149	0.0165	0.0181
80	0.0122	0.0139	0.0158
120	0.0132	0.0158	0.0184

Table 19: Degraded Link Change proportion Results for Remote Subscription with Home Tunneling

Number of Mobile Nodes	90 % Confidence Interval Lower Bound	Mean	90 % Confidence Interval Upper Bound
40	0.0010	0.0015	0.0020
80	0.0006	0.0013	0.0020
120	0.0004	0.0014	0.0025

Table 20: Degraded Link Change proportion Results for Remote Subscription with Minimal Multicast Encapsulation

Number of Mobile Nodes	90 % Confidence Interval Lower Bound	Mean	90 % Confidence Interval Upper Bound
40	0.0008	0.0012	0.0016
80	0.0012	0.0014	0.0015
120	0.0003	0.0010	0.0018

4.5.2.2 Packet Loss Proportions

This subsection presents the proportion of packet lost due to degraded link changes in tables 20 through 23. These values tend to be much smaller than the degrade link change proportion simply because there are many more packets that can be received than there are link changes (roughly 7000 link changes versus more than 870,000 packets received), and in most cases there tended to be about one packet lost per degraded link. This rate should increase as the transmission rate increases simply because more packets will be transmitted to before the home agent receives an update or before the mobile node successfully joins the multicast tree.

Table 21: Packet loss proportion results for Bi-directional tunneling with Home Tunneling

Number of Mobile Nodes	90 % Confidence Interval Lower Bound	Mean	90 % Confidence Interval Upper Bound
40	1.28E-05	1.46E-05	1.64E-05
80	1.14E-05	1.53E-05	1.91E-05
120	0	4.53E-04	1.49E-04

Table 22: Packet loss proportion results for Bi-directional Tunneling with Minimal Multicast Encapsulation

Number of Mobile Nodes	90 % Confidence Interval Lower Bound	Mean	90 % Confidence Interval Upper Bound
40	1.4E-05	1.7E-05	2E-05
80	1.2E-05	1.6E-05	2.1E-05
120	0	4.5E-4	1.5E-3

Table 23: Packet loss proportion results for Remote Subscription with Home Tunneling

Number of Mobile Nodes	90 % Confidence Interval Lower Bound	Mean	90 % Confidence Interval Upper Bound
40	8.3E-07	1.3E-06	1.7E-06
80	5.1E-07	1.1E-06	1.8E-06
120	3.5E-7	1.36E-6	2.37E-06

Table 24: Packet loss proportion results for Remote Subscription with Minimal Multicast Encapsulation

Number of Mobile Nodes	90 % Confidence Interval Lower Bound	Mean	90 % Confidence Interval Upper Bound
40	7.4E-07	1E-06	1.3E-06
80	1E-06	1.2E-06	1.4E-06
120	2.7E-07	9.7E-07	1.7E-06

With two notable exceptions, the combinations that employ bi-directional tunneling follow the same general trend of being an order of magnitude greater than those that employ remote subscription. The two exceptions are for the bi-directional tunnel 120 node scenarios that suffer an even greater amount of packet loss (see Tables 20 and 21). This is due to the fact that, for one seeds, there was an inordinately high amount of packet

loss experienced for the 120 node scenario runs of both the bi-directional tunnel combination. These results are presented in Table 24. Upon further analysis, it was discovered that one mobile node failed to join the multicast group after its first link change, and that this failure to join was responsible for the increase in lost packet. This node, however, functioned properly on all subsequent link changes. While the exact cause of this a failure could not be determined, it is noted that if the packet lost due to this link failure are discounted, results similar to the other three seed runs are obtained.

Table 25: Anomalous results for 120 Node Scenarios using Bi-directional Tunnel Combinations

Mechanism Combination	Seed	Degraded Links	Packets Lost due to Link Change
Bi-directional Tunnel with Home Tunnel	128	105	112
Bi-directional Tunnel with Home Tunnel	371	104	116
Bi-directional Tunnel with Home Tunnel	754	116	128
Bi-directional Tunnel with Home Tunnel	539257	114	14911
Bi-directional Tunnel with Minimal Multicast Encapsulation	128	106	113
Bi-directional Tunnel with Minimal Multicast Encapsulation	371	111	117
Bi-directional Tunnel with Minimal Multicast Encapsulation	754	138	147
Bi-directional Tunnel with Minimal Multicast Encapsulation	539257	127	13110

4.5.3 Analysis of Required Throughput

As previously stated, the required throughput measures transmission load placed on a mobility agent while providing mobility support to mobile nodes. The throughput for the

most heavily loaded node for each transmit-receive combination is presented in Table 25.

The increased load experienced by the bi-directional tunnel combinations is attributed to multiple tunnels converging at one foreign agent. The slightly higher values for bi-directional tunneling with minimal multicast encapsulation can be attributed to the additional overhead imposed by encapsulating the original multicast packet.

Table 26: Maximum Observed Required Mobility Agent Throughput

Mechanism Combination	Number of Receivers	Maximum Base Station Radio Link Throughput (bits per second)
Bi-directional Tunnel with Home Tunnel	120	53,743,850
Bi-directional Tunnel with Minimal Multicast Encapsulation	120	57,850,209
Remote Subscription with Home Tunnel	120	2,721,119
Remote Subscription with Minimal Multicast Encapsulation	120	2,976,071
Bi-directional Tunnel with Home Tunnel	80	10,397,252
Bi-directional Tunnel with Minimal Multicast Encapsulation	80	11,189,800
Remote Subscription with Home Tunnel	80	1,202,742
Remote Subscription with Minimal Multicast Encapsulation	80	1,315,864
Bi-directional Tunnel with Home Tunnel	40	5,313,128
Bi-directional Tunnel with Minimal Multicast Encapsulation	40	5,718,731
Remote Subscription with Home Tunnel	40	860,007
Remote Subscription with Minimal Multicast Encapsulation	40	940,995

4.6 Summary

This chapter began with a presentation of statistical methods used in analysis of scenario results. Next simulation scenarios were discussed along with methodologies for the collection of data. Finally, results were presented, analyzed and compared for each performance metric: hop and distance ratios, degraded link and packet loss proportions, and required throughput.

Chapter 5: Conclusions and Recommendations

5.1 Restatement of Research Goal

This research had two goals:

- to compare the performance of currently proposed IP mobility support mechanisms for IP multicast in terms of routing efficiency, packet loss and mobility agent loading.
- to introduce and determine the feasibility of a novel mobile multicast transmission support mechanism.

5.2 Conclusions

5.2.1 Results Synopsis

Four combinations of mobile IP multicast transmission were compared. The two currently proposed IETF mobile IP multicast reception mechanisms, bi-directional tunneling and remote subscription, were paired with the IETF home tunneling transmission mechanism as well as with the minimal multicast encapsulation mechanism introduced in this research. Three areas of performance were examined in this research. These areas were routing efficiency, packet loss, and the required mobility agent throughput.

Results indicate that the combination of remote subscription and minimal multicast encapsulation gave the best routing efficiency—path ratios of 1—while unicast tunneling, either for transmission or reception greatly decreases routing efficiency. The worst-case

routing efficiencies were suffered when unicast tunneling was used for both sending and receiving.

In the area of packet loss, any combination that utilized bi-directional tunneling suffered much greater packet losses due to link changes than did combinations that did not utilize bi-directional tunneling. This is because the inherent delay associated with updating the home agent is greater than the delay required to join the multicast tree as would be required in remote subscription.

With regard to required mobility agent throughput, this research shows that tunnel convergence caused by bi-directional tunneling can increase loading on mobility agents by a factor of almost 20. Finally, minimal multicast encapsulation was shown to slightly increase the required throughput because of encapsulation.

5.2.2 Recommendations

For improved path efficiency, it is recommended to avoid tunneling for mobility support to IP multicasting. Remote subscription provides the best path efficiency of the tested reception mechanisms, but as noted in [ChW98] it does not support locally scoped multicast groups. If access to locally scoped groups is required, a tunneling mechanism such as bi-directional tunneling or MoM should be used. A mechanism such as MoM should be chosen over bi-directional tunneling because MoM attempts to eliminate tunnel convergence.

Minimal multicast encapsulation is recommended for improving transmission path efficiency when IP addresses are limited and when proper return IP addresses are required. Otherwise, if bandwidth is limited, IP addresses are plentiful, and return

addresses are not required in the IP packet, a transmission mechanism utilizing co-located care-of addresses is recommended, to conserve overhead required for minimal multicast encapsulation. Also, in bandwidth-limited networks, bi-directional tunneling is not recommended because of greatly increased load due to tunnel convergence.

5.3 Significant Results of Research

This work is the first to model the operation of IP mobility mechanisms for IP multicast at the IP protocol level. This work also introduces and verifies the viability the novel mobile IP multicast transmission mechanism, minimal multicast encapsulation. Finally this work provides a model that can be used to test a variety of aspects of mobile IP multicasting.

5.4 Future Research

The models used in this research can be used to further study the effect of varying traffic levels on IP mobility support protocols performance. They can also be used to examine the interaction between IP mobility support for multicast and IP mobility support for unicast. The effects of increasing the number of mobile sources, as well as interaction between mobile and fixed sources and receivers can also be examined. With modifications the models could be expanded to simulate and analyze the operation of the MoM multicast support mechanism, as well as any new mobile multicast support mechanisms that may arise.

Finally the flexibility of the models and of OPNET, the models could be used to study the effects of disparate bandwidth capabilities on different subnets, a real world

concerned raised in [BrS96]. These models could be modified and used to analyze and compare solutions for allocating available bandwidth to newly arriving nodes.

Bibliography

- [Bal97] A. Ballardie, "Core Based Trees (CBT Version 2) Multicast Routing," RFC 2189, Internet Engineering Task Force, September 1997.
- [BrS96] K. Brown, S. Singh, "The Problem of Multicast in Mobile Networks", IEEE 5th International Conference on Computer Communications and Networks, October 1996.
- [ChW98] V. Chickarmane, C. Williamson, R. Bunt, W. Mackrell "Multicast support for mobile hosts using Mobile IP: Design issues and proposed architecture," *Mobile Networks and Applications*, Vol. 3 No.4, Baltzer Science Publications, Bussum, 1998.
- [DeC90] S. Deering and D. Cheriton, "Multicast Routing in Datagram Internetworks and Extended LANs," *ACM Transactions on Computer Systems*, Vol. 8 No. 2, May 1990, Pages 85-110.
- [Dee89] S. Deering, "Host Extensions for IP Multicasting" RFC 1112, Internet Engineering Task Force, August 1989.
- [Dee99] S. Deering, et al, "Protocol Independent Multicast Version 2 Dense Mode Specification" Internet Draft: draft-ietf-pim-v2-dm-03.txt, Internet Engineering Task Force, March 1999.
- [Fen97] W. Fenner, "Internet Group Management Protocol, Version 2," RFC 2236, Internet Engineering Task Force, November 1997.
- [FID99] A. Fladenmuller and R. De Silva, "The effect of Mobile IP handoffs on the performance of TCP," *Mobile Networks and Applications*, Vol. 4, 1999, pages 131-135.
- [GoN99] M. Goncalves and K. Niles, *IP Multicasting: Concepts and Applications*, McGraw-Hill, New York 1999.
- [HaW97] T. Harrison, C. Williams, W. Mackrell, and R. Bunt, "Mobile multicast (MoM) protocol: multicast support for mobile hosts," *MOBICOM '97. Proceedings of the third annual ACM/IEEE international conference on computing and networking*, 1997, pages 151-160.
- [LaK00] A. Law and W. Kelton, *Simulation Modeling and Analysis*, McGraw Hill, Boston 2000.
- [Mau98] T. Maufer, *Deploying IP Multicast in the Enterprise*, Prentice-Hall, Upper Saddle River, 1998.
- [Mck99] K. McKay, "RTP Payload Format for PureVoice(tm) Audio," RFC 2658, Internet Engineering Taskforce, August 1999.
- [Mil99] K. Miller, *Multicast Networking and Application*, Addison-Wesley, Reading, 1999.

- [MiA95] J. Milton and J Arnold, *Introduction to Probability and Statistics: Principles and Applications for Engineering and the Computing Sciences*, Irwin McGraw-Hill, Boston, 1995.
- [Per96a] C. Perkins, "IP Mobility Support", RFC 2002, Internet Engineering Task Force, October 1996.
- [Per96b] C. Perkins, "IP Encapsulation within IP", RFC 2003, Internet Engineering Task Force, October 1996.
- [Per96c] C. Perkins, "Minimal Encapsulation within IP", RFC 2004, Internet Engineering Task Force October 1996.
- [Wen98] L. Wentz, *Lessons from Bosnia: The IFOR Experience*. Washington, DC: National Defense University Press, April 1998.
- [WiH98] C. Williamson, T. Harrison, W. Mackrell, R. Bunt, "Performance Evaluation of the MoM mobile multicast protocol," *Mobile Networks and Applications*, Vol. 3 No.2, Baltzer Science Publications, Bussum, 1998.

Vita

Captain Muller was born on 26 December 1965 in Auburn, Indiana. He served in the United States Marine Corps from 1987 to 1995. In 1994, he received a Bachelor of Science in computer science from the University of Missouri. He was commissioned a Second Lieutenant in the United States Air Force after completing of Officer Training School in 1996. He served as a software evaluator on the Cheyenne Mountain Upgrade test team at Detachment 4, Air Force Operational Test and Evaluation Center at Peterson Air Force base Colorado. He has earned the Air Force Commendation Medal and the Navy Achievement Medal.

He is married to the former Hyon-I Son of Moon-san, Republic of Korea. They have two delightful children: Emerald and Joshua.

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE April 2000		3. REPORT TYPE AND DATES COVERED Master's Thesis
4. TITLE AND SUBTITLE A COMPARATIVE ANALYSIS OF PROPOSED MOBILITY SUPPORT SCHEMES FOR IP MULTICAST			5. FUNDING NUMBERS	
6. AUTHOR(S) Alexander Muller Jr., Captain, USAF				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Air Force Institute of Technology 2950 P Street Wright-Patterson AFB OH 45433			8. PERFORMING ORGANIZATION REPORT NUMBER AFIT/GCS/ENG/00J-02	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Communication Agency/ITAI Ken Fore 203 W. Losey St Scott Air Force Base IL 62225			10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES Richard A. Raines, Major, USAF DSN: 785-3636, ext. 4715				
12a. DISTRIBUTION AVAILABILITY STATEMENT Approved for public release; distribution unlimited			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) This thesis introduces a novel mobile multicast transmission mechanism called Minimal Multicast Encapsulation. Additionally, this thesis analyzes the performance of mobility support schemes for IP multicast. Specifically, it compares the performance of combinations of two receive mechanisms and two transmit mechanisms. The receive mechanisms are the Internet Engineering Task Force (IETF) bi-directional tunneling mechanism and the IETF remote subscription mechanism. The transmit mechanisms are the IETF home tunneling mechanism and the Minimal Multicast Encapsulation. The performance analysis consists of examining path efficiencies, packet loss rates, and required mobility agent throughputs for each of the four possible combinations of the abovementioned transmit and receive mechanisms. Results of the analysis indicate that combinations that include bi-directional tunneling or home tunneling suffer from average path lengths at least 2 times the optimal path length. The combination of bi-directional tunneling and home tunneling has average path lengths that are 3 to 5 times optimal. The combination of remote subscription and Minimal Multicast Encapsulation provided optimal path lengths. Bi-directional tunneling suffered from roughly 10 times more degraded link changes due to packet loss than remote subscription. Maximum throughput requirements were 20 times greater for bi-directional tunneling than the maximum throughput requirements for remote subscription.				
14. SUBJECT TERMS IP Mobility, Minimal Multicast Encapsulation, Mobile IP, Mobile Multicast, Mobile Networks, Multicast			15. NUMBER OF PAGES 113	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UL	